Decomposing Shifts in the Beveridge Curve: Implications for Labor Market Dynamics and Inflation*

by

Katharine Abraham
University of Maryland
and NBER

John Haltiwanger
University of Maryland
and NBER

Lea Rendell
Kenyon College

March 5, 2026

*This is a draft of a paper to be presented at the March 12-13, 2026 NBER "Inflation: Frontiers of Research and Policy" conference. We thank Martim Leitao for excellent research assistance. We also thank Regis Barnichon and Adam Shapiro for sharing MSA level data with us.

I. **Introduction**

In the traditional Phillips curve equation, the unemployment rate serves as the measure of economic conditions. Its inclusion is motivated by the idea that, when unemployment is low, employers may find it harder to fill available jobs, putting upward pressure on wages and thus on prices. Employers' recruiting success, however, depends on what is happening not only on the supply side of the labor market but also on the demand side of the market. Several recent papers have adopted the ratio of vacancies to unemployment, a measure incorporating information on both unmet labor demand and unutilized labor supply, in place of the unemployment rate as the driving variable for Phillips curve estimation (Ball, Leigh and Mishra 2022; Barnichon and Shapiro 2024; Benigno and Eggertsson 2024). If the relationship between vacancies and unemployment—the Beveridge curve—is stable, the vacancy-to-unemployment ratio should move together with the unemployment rate, and either should serve equally well for Phillips curve estimation. Because the Beveridge curve is not stable, however, one might expect the ratio of vacancies to unemployment or some related measure to be more informative than the unemployment rate about employers' hiring experiences and thus the likelihood of upward pressure on wages and prices.

As a measure of labor market tightness, the vacancy-to-unemployment ratio has the virtue of incorporating information about the demand side of the labor market, but like the unemployment rate, it treats the potential supply of workers available to fill employers' vacant jobs as just the number of people who are unemployed. Even among the unemployed, however, some people may add more to the effective pool of available workers than others; the long-term unemployed, for example, are significantly less likely to be hired from one month to the next than those who only recently became unemployed. Just as important, employers commonly

recruit new hires not only from among the unemployed but also from among people who are out of the labor force or currently employed elsewhere. Several recent papers have developed alternative measures of the available worker pool that build on this insight (e.g., Hornstein, Kudlyak and Lange 2014; Hall and Shulhofer-Wohl 2018; Abraham, Haltiwanger and Rendell 2020).

The present paper follows the approach taken by Abraham, Haltiwanger, and Rendell (2020) (hereafter AHR2020) for constructing a measure of effective searchers that accounts for the possibility of hiring from different groups among the unemployed, those out of the labor force and the currently employed. The measure of effective searchers is constructed by taking a weighted sum across twenty-two different groups within the population age sixteen and older, with the weights based on the relative base period job finding rates for each of the different groups as a proxy for relative job search intensities. Measured in this way, relative job search intensities are especially high among the recently unemployed but non-zero for all groups. AHR2020 propose a generalized matching function that accounts for the availability of potential hires from all of these sources and show that this generalized matching function better captures actual job-filling and job-finding dynamics. They also note—but do not fully explore—the potential implications of their generalized analysis for interpreting shifts in the standard Beveridge Curve and for understanding wage and price pressures arising in the labor market.

Shifts in the standard Beveridge curve commonly are characterized as indicative of changes in matching efficiency (e.g., Barlevy et al. 2024). Because unemployment is an imperfect proxy for the labor supply potentially available to employers, however, conclusions about matching efficiency based on the standard analysis may be misleading. While acknowledging that effective searchers also is an imperfect measure of the potentially available

labor supply, we view the measure of matching efficiency derived from the generalized model as a substantial improvement over the standard measure. Based on an analysis that relates the standard matching function to the generalized matching function, we show that shifts in the standard Beveridge curve reflect not only changes in the efficiency with which effective searchers are matched to vacant jobs ("true" generalized matching efficiency) but also changes in the ratio of effective searchers to unemployment, along with an additional term that reflects other factors.

Next, we derive approximate decompositions that allow us to analyze movements in the vacancy–unemployment ratio, the standard measure of labor market tightness. Starting with the standard approach, we show that changes in this ratio can be separated into three components—movements along the standard Beveridge curve; outward or inward shifts in the curve attributed to changes in matching efficiency as conventionally measured; and shifts driven by other factors, such as changes in the separation rate. While consistent with the standard framework, this decomposition—based on matching efficiency measured in the usual way—has important shortcomings in accounting for observed variation in labor market tightness. The problem stems from an inconsistency between the matching efficiency implied by the standard matching function and the matching efficiency consistent with the job-finding rate of the unemployed.

It also is possible to use the generalized framework, which builds on a matching function specified in terms of vacancies and effective searchers, to analyze movements in the vacancy-to-unemployment ratio. In addition to terms analogous to those in the conventional approach, the generalized decomposition includes a component that captures shifts driven by changes in the ratio of effective unemployed searchers to the number of unemployed people. This approximation tracks observed standard labor market tightness closely. Within the generalized

model, the matching efficiency derived from the generalized matching function aligns closely with the matching efficiency implied by the job-finding rate of the unemployed. Taken together, these results provide further support for the generalized framework.

Building on these findings, we next decompose generalized labor market tightness into its underlying drivers. Changes in generalized tightness reflect movements along the generalized Beveridge curve, shifts in that curve due to changes in generalized matching efficiency, and shifts due to other factors. This decomposition is particularly informative for understanding the relationship between generalized labor market tightness and inflation. The results are consistent with the view that generalized tightness provides a more comprehensive measure of labor market conditions.

Our empirical analysis spans the period from 1994:Q1 to 2025:Q2, covering key episodes such as the Great Financial Crisis, the COVID-19 pandemic, and the subsequent pandemic recovery. We show that, over this period, standard matching efficiency behaves quite differently than generalized matching efficiency. For example, during the Great Financial Crisis, matching efficiency measured in the usual way fell markedly, but generalized matching efficiency fell much less. Much of the apparent decline in matching efficiency measured in the standard way is instead explained by the fact that effective searchers fell relative to the number of unemployed people. As a result, the increase in the availability of potential workers was proportionally smaller than the increase in unemployment. In the post-pandemic period, to take another example, standard matching efficiency largely rebounded from its pandemic low, but generalized matching efficiency continued to drift downwards.

Our empirical analysis of labor market tightness focuses primarily on the generalized tightness measure (the ratio of vacancies to effective searchers). Much of the movement in

generalized tightness reflects movement along the generalized Beveridge curve, but shifts due to changes in generalized matching efficiency and shifts for other reasons also have played important roles. During the Great Financial Crisis, for example, increases in the number of effective searchers reflecting movements along the generalized Beveridge curve and Beveridge curve shifters other than matching efficiency both worked to reduce the ratio of vacancies to effective searchers, whereas reductions in generalized matching efficiency that on their own would have shifted the Beveridge curve outwards worked in the opposite direction.

With this decomposition of generalized labor market tightness in hand, we turn to the relationship between generalized tightness and inflation dynamics. Our analysis builds on recent work by Ball, Leigh and Mishra (2022), Barnichon and Shapiro (2024) and Benigno and Eggertsson (2024), all of whom study the relationship between labor market tightness measured by the vacancies-to-unemployment ratio and inflation. Ball, Leigh and Mishra (2022) motivate and document a nonlinear relationship between standard tightness and inflation. Similarly, Benigno and Eggertsson (2024) make the case that, once labor markets become sufficiently tight, the Phillips curve becomes essentially vertical. Barnichon and Shapiro (2024) use less formal decompositions of standard tightness to argue that shifts in the Beveridge curve—particularly shifts in matching efficiency—play a central role in explaining inflation dynamics.

Similar to other results reported by Barnichon and Shapiro (2024), we find that generalized tightness outperforms standard tightness in a simple Phillips curve equation. Models that allow for nonlinear tightness effects do better than models in which tightness enters linearly, but there too, the generalized model outperforms the standard model. Applying the decomposition of generalized tightness described above, we also estimate a model in which the terms of the decomposition enter the Phillips curve equation separately. In these models, changes

in the vacancy-to-effective searcher ratio due to movements along the generalized Beveridge curve are allowed to affect inflation differently than changes in that ratio due to changes in generalized matching efficiency or to changes in other factors that shift the generalized Beveridge curve. These results suggest that wages and thus prices may respond more to changes in tightness that employers view as likely to be lasting as opposed to changes that are likely to be reversed over the course of the business cycle.

The framework we develop here complements and extends earlier analyses in several important respects. First, based on an exact decomposition, we analyze the relationship between the standard matching function and its generalized counterpart. Our empirical analysis of this relationship reveals that much of what has been characterized as changes over time in matching efficiency within the standard Beveridge curve framework instead reflects changes in the ratio of effective searchers to unemployment. Second, we analyze the relationship between standard and generalized labor market tightness and provide decompositions that illuminate the forces underlying their movements. Third, in addition to comparing the performance of standard and generalized measures of labor market tightness in empirical Phillips curves, we show both that there appear to be nonlinearities in these measures' effects and that the effect on inflation can differ depending on the reason for the changes in the generalized measure. The latter finding suggests that inflation responds not only to the level of tightness in the labor market but also to its causes, with changes in tightness due to drivers that exhibit more persistence having larger effects than changes from other sources.

II. Standard and Generalized Labor Market Tightness

The standard matching function—which links hires, job-filling rates, and job-finding rates to vacancies and unemployment—serves as a central workhorse of the macroeconomic

labor literature in both theoretical and empirical applications. We begin by reviewing this benchmark framework and then highlight an important limitation—namely, that it largely neglects, both conceptually and empirically, the substantial share of hires that occur as movements from out of the labor force into employment or job-to-job transitions. To address this gap, we introduce a generalized matching function that incorporates these additional flows. We then examine the implications of this broader framework for measuring labor market tightness and for understanding the forces that drive changes in these measures.

The Standard and Generalized Matching Functions

In the standard specification of the canonical search-and-matching model (Diamond 1982; Blanchard and Diamond 1989, 1992; Mortensen and Pissarides 1994; Pissarides 2000), employers create job openings they would like to fill, and unemployed individuals search among these job openings for employment. The process of matching unemployed workers to vacant jobs is represented by a production function, often assumed to be Cobb-Douglas in form, with vacancies and unemployment as the inputs and matches (hires) as the output:

$$(1) \quad h_t = m(v_t, u_t) = \mu_t^s v_t^{1-\alpha_s} u_t^{\alpha_s}$$

where h_t is the hiring rate, v_t is the job opening rate, u_t is the unemployment rate, t is the time period, μ_t is a time-varying match efficiency parameter, and $\alpha_s(1-\alpha_s)$ is the elasticity of the (standard) matching function with respect to unemployment (vacancies).[1] In this framework, labor market tightness is expressed as:

[1]So long as it is constant returns to scale, the matching function can be written either in levels or in rates relative to a common denominator. We write it in rates. In our empirical analysis, the common denominator is the population aged 16 years and older.

(2) $$\theta_t = \frac{v_t}{u_t}.$$

In the standard model steady state, the number of separations (inflows to unemployment) must equal the number of hires (outflows from unemployment). This steady state relationship can be expressed:

(3) $$\delta_t = h_t = m\left(v_t, u_t\right) = \mu_t^s v_t^{1-\alpha_s} u_t^{\alpha_s},$$

where δ_t is the separation rate. The downward sloping relationship between the vacancy rate and the unemployment rate implied by equation (3) commonly is termed the Beveridge curve. In this framework, shifts in either δ_t or μ_t^s will shift the position of the Beveridge curve. Improvement in the matching function (an increase in μ_t^s), for example, shifts the Beveridge curve inward (lowering vacancies for given unemployment), while deterioration in the matching function shifts the Beveridge curve outward (raising vacancies for given unemployment). In addition, out-of-steady-state dynamics lead to shifts in the Beveridge curve.

In the model as just sketched out, $\theta_t = \frac{v_t}{u_t}$, but unemployment and vacancies are imperfect proxies for the measures of effective searchers and effective vacancies that we believe should be the objects of interest. As discussed in AHR2020, the standard framework neglects the substantial fraction of hires who are individuals coming directly from out of the labor force or from another job . Consistent with the recent literature, AHR2020 show that the standard matching function requires substantial variation in matching efficiency to account for observed shifts in the Beveridge curve, as well as to explain the considerable residual variation between actual job-filling and job-finding rates and those implied by the model. An important reason for this is that ignoring the broader stock of potential job candidates as well heterogeneity in the job

search intensities of specific groups yields a mis-specified matching function. The standard framework also ignores the possibility of variation over time in employer search intensity.

Following AHR2020, we can elaborate the simple model to account for the first of these complexities. In this paper, we focus on the importance of more accurately measuring the availability of potential hires by accounting for different groups among the unemployed and also hires that may come from among those currently out of the labor force or in another job. Accounting for changes over time in employer recruiting intensity, as discussed in Davis, Faberman and Haltiwanger (2013), is left for future work.[2]

Building on the standard hiring function, we can write:

$$(4) \qquad h_t = m\left(v_t, \sum_i \rho_t^{l_i} l_{it}\right) = \mu_t^g v_t^{1-\alpha_g} \left(\sum_i \rho_t^{l_i} l_{it}\right)^{\alpha_g},$$

where l_i represents job searchers of type i and $\rho_t^{l_i}$ represents the intensity of job search on the part of searchers of type i at time t. Note that the elasticity of the matching function specified in equation (4) as well as the unexplained residual variation in the actual number of matches as compared to the number implied by the model are likely to differ from those associated with the standard model specified in equation (1). This underlies our use of the superscript g. In practice, our measures of job search intensity are job-finding rates in a chosen base period relative to the job-finding rate of a reference group in that same period.

In this expanded framework, labor market tightness can be written as:

[2] Davis, Faberman and Haltiwanger (2013) provide evidence of establishment-level variation in recruiting intensity, estimating that the elasticity of the job filling rate with respect to the gross hires rate at the establishment level is 0.82. Assuming that this micro elasticity holds at the aggregate level, they postulate a measure of aggregate recruiting intensity equal to $h_t^{0.82}$. While their work provides reason to think that aggregate recruiting intensity is likely to vary, the reliance on variation in aggregate hires to infer that variation is not well suited for the decomposition approaches we consider here. Recognizing that unmeasured variation in recruiting intensity is part of measured matching efficiency, we see the development and incorporation of alternative measures of recruiting intensity as a priority for future work.

$$(5) \qquad \tilde{\theta}_t = \frac{v_t}{\sum_i \rho_t^{l_i} l_{it}}.$$

We refer to the denominator of equation (5) as effective searchers (ES_t). The steady state equilibrium of hires equal to separations is now:

$$(6) \qquad \delta_t = h_t = m\left(v_t, \sum_i \rho_t^{l_i} l_{it}\right) = \mu_t^g \left(v_t\right)^{1-\alpha_g} \left(\sum_i \rho_t^{l_i} l_{it}\right)^{\alpha_g},$$

where separations are now all separations from employment and all variables are again expressed as rates. Over the course of a business cycle, absent changes in matching efficiency or other factors that shift the position of the generalized Beveridge curve, effective searchers and effective vacancies will move inversely as implied by equation (6).

Decomposing Standard Matching Efficiency

With these two alternative frameworks, consider log standard matching efficiency and log generalized matching efficiency:

$$(7) \quad \log(\mu_t^s) = \log(h_t) - (1-\alpha^s)\log(v_t) - \alpha^s \log(u_t)$$

$$(8) \quad \log(\mu_t^g) = \log(h_t) - (1-\alpha^g)\log(v_t) - \alpha^g \log(ES_t)$$

where $ES_t = \sum_i \rho_t^{l_i} l_{it}$. Combining these two approaches yields:

$$(9) \quad \log(\mu_t^s) = \log(\mu_t^g) + \alpha^g \log(ES_t / u_t) + (\alpha^s - \alpha^g)\log(\theta_t)$$

In words, changes in standard matching efficiency reflect three things: 1) changes in (properly specified) generalized matching efficiency; 2) the ratio of effective searchers to unemployment; and 3) to the extent that the empirical elasticity of the matching function differs between the standard the generalized specification, movements in standard tightness. The results in AHR2020 already provide *prima facie* evidence that this decomposition is likely to be informative. More

specifically, their generalized matching function tracks actual job filling and finding rates more closely than the standard matching function. Put differently, the standard matching function requires greater implied changes in matching efficiency to match observed job filling and job finding rates.

In the empirical analysis reported later in the paper, we apply this decomposition to shed light on the distinction between standard and generalized matching efficiency and the interpretation of shifts in the standard Beveridge curve. The decomposition also is relevant for understanding the relationship between inflation and indicators of labor market tightness.

Job Finding Rates from Unemployment in the Standard versus the Generalized Model

To draw out the differences between the standard and generalized model, it is helpful to characterize what each implies about the job finding rate for the unemployed. Note that, in the general model, since hires may come from out of the labor force or through job-to-job transitions, the ratio h_t/u_t is not the job finding rate for the unemployed. To draw out the implications of the generalized model for the rate of hiring from unemployment, we build on insights from AHR2020 which in turn built on Hall and Schulhofer-Wohl (2018). In the general model, job-finding rates differ across groups only because of differences in their effective search intensity, meaning that they are the same on a search intensity–adjusted basis. The common search-intensity-adjusted job-finding rate varies with the tightness of the labor market:

$$\frac{h_t}{ES_t} = \frac{h_{it}}{ES_{it}} = \frac{h_{it}}{\rho_t^{l_i} l_{it}} = \tilde{f}_t = \tilde{f}_{it} = A_t^g T_t^{\eta_g}, \tag{10}$$

where A_t is a transformation of matching efficiency such that $A_t^g = (\mu_t^g)^{1+\eta_g}$, $T_t = v_t / h_t$ is average vacancy duration and $\eta^g = (1-\alpha^g)/\alpha^g$).[3]

Using this approach, the job-finding rate for the unemployed can be written as:

$$(11) \qquad \frac{ue_t}{u_t} = (\mu_t^g)^{1+\eta^g} (v_t / h_t)^{\eta^g} \sum_{i \in u} \left(\rho_t^{l_i} l_{it} / l_{ut} \right).$$

where ue_t is the flow from unemployment to employment (i.e., hires from unemployment). An attractive feature of this approach is that the standard model is nested naturally as:

$$(12) \qquad \frac{ue_t}{u_t} = (\mu_t^s)^{1+\eta^s} \left(\frac{v_t}{h_t} \right)^{\eta^s}$$

Equations (11) and (12) enable traditional labor market tightness (*v/u*) to be decomposed in two distinct ways. We combine each of these equations in turn with the flow equation for unemployment:

$$(13) \quad \Delta u_{t+1} = eu_t - ue_t + nu_t - un_t$$

where eu_t is the flow from employment to unemployment, nu_t is the flow from not in the labor force to unemployment, and un_t is the flow from unemployment to not in the labor force.

Combining (12) with (13) and taking a first order log approximation yields (see Appendix C for details):

$$(14) \quad \hat{v}_t - \hat{u}_t = -\left(\frac{\alpha^s}{1-\alpha^s} + 1\right)\hat{u}_t - \frac{1}{1-\alpha^s}\hat{\mu}_t^s + \widehat{x_t^s}$$

where $\hat{y}_t = \ln y_t - \ln \bar{y}$, $\bar{y}$ is the mean of *y*, and we assume $\Delta u_{t+1} = 0$.[4] The right-hand side of equation (14) can be summarized as three composite terms: 1) a term due to *u* reflecting

[3] AHR2020 show that equation (10) is an exact transformation of equation (4), the generalized matching function.
[4] Ahn and Crane (2020) and Barnichon and Shapiro (2024) carry out similar decompositions.

movement along the Beveridge Curve given by $-\left(\frac{\alpha^s}{1-\alpha^s}+1\right)\hat{u}_t$; 2) a term reflecting Beveridge curve shifts due to changes in matching efficiency given by $-\frac{1}{1-\alpha^s}\hat{\mu}_t^s$; and 3) a term reflecting shifts in the Beveridge curve from factors other than matching efficiency given by $\widehat{x_t^s}=\frac{\alpha^s}{1-\alpha^s}\left[\frac{\bar{e}u}{\bar{\phi}}\widehat{eu}_t+\frac{\bar{n}u}{\bar{\phi}}\widehat{nu}_t-\frac{\bar{u}n}{\bar{\phi}}\widehat{un}_t-\frac{\Delta u_{t+1}}{\bar{\phi}}\right]+\hat{h}_t$ (with $\bar{\phi}=\bar{e}u+\bar{n}u-\bar{u}n$). Following similar steps, (11) and (13) can be combined and approximated log linearly to yield:

$$\hat{v}_t-\hat{u}_t=-\left(\frac{\alpha^g}{1-\alpha^g}+1\right)\hat{u}_t-\frac{1}{1-\alpha^g}\hat{\mu}_t^g+\widehat{x_t^g}-\frac{\alpha^g}{1-\alpha^g}\hat{Z}_t \quad (15)$$

where $Z_t=\frac{\sum_{i\in u}\rho^{l_i}l_{it}}{u_t}$ and $\widehat{x_t^g}=\frac{\alpha^g}{1-\alpha^g}\left[\frac{\bar{e}u}{\bar{\phi}}\widehat{eu}_t+\frac{\bar{n}u}{\bar{\phi}}\widehat{nu}_t-\frac{\bar{u}n}{\bar{\phi}}\widehat{un}_t-\frac{\Delta u_{t+1}}{\bar{\phi}}\right]+\hat{h}_t$ Equation (15) has a structure similar to equation (14) but with α^g in place of α^s and an additional term reflecting the ratio of effective searchers among the unemployed to the number of people unemployed. Thus, interestingly, the generalized model has implications for the evolution of standard labor market tightness.

In implementing equations (14) and (15), we use the matching efficiency that is consistent with equations (7) and (8), that is, the matching efficiency that makes hires consistent with the levels of job openings and either unemployment from the standard matching function or effective searchers from the generalized matching functions. As a specification check, we also can construct the matching efficiency that is consistent with equations (11) and (12), the matching efficiency consistent with the standard and generalized *ue/u* flow equations.

Decomposing Generalized Tightness

Generalized tightness can be decomposed following a similar approach. To do this, we combine the generalized matching function of equation (4) with two additional relationships: 1)

$e_t + u_t + n_t = 1$ where the three terms are employment, unemployment and not in the labor force as fractions of the working age population and 2) $\Delta e_{t+1} = h_t - sep_t$ where h_t is hires during period t and sep_t is all separations from employment in period t. Combining and log linearizing following steps similar to those above yields the following decomposition for the generalized Beveridge curve:

$$(16) \qquad \widehat{v_t} - \widehat{\mathrm{ES}_t} = -\left(\frac{\alpha^g}{1-\alpha^g} + 1\right)\widehat{ES_t} - \frac{1}{1-\alpha^g}\hat{\mu}_t^g + \hat{x}_t^g$$

where $\widehat{x_t} = [\widehat{sep_t} - \frac{\Delta u_{t+1}}{xss} - \frac{\Delta n_{t+1}}{xss}]$, $xss = \overline{sep}$ (assuming in steady state (on average) that Δu_{t+1} and Δn_{t+1} are equal to zero), and $ES_t = \sum_i \rho_t^{s_i} s_{it}$. This decomposition has a similar structure to that for standard labor market tightness but movement along the generalized Beveridge curve involves effective searchers. In implementation, we use the matching efficiency that is consistent with equation (8), that is, the matching efficiency consistent with total hires.

III. Data and Methods

Our empirical analysis begins with an examination of Beveridge curves and measures of labor market tightness using national-level data. We construct these measures as described in AHR2020. The standard Beveridge curve is the downward-sloping relationship between vacancies and unemployment; our implementation uses v_t and u_t, vacancy and unemployment rates expressed as a share of the population aged 16 plus. Standard labor market tightness is v_t/u_t. The generalized Beveridge curve plots the relationship between v_t and ES_t, where ES_t is effective searchers, defined as $\sum_i \rho_t^{l_i} l_{it}$, with $\rho_t^{l_i}$ the relative job-finding rate for group i in period t and l_{it}

the group i population share. We interpret cross-group variation in job-finding rates as variation in search intensity.[5] Generalized labor market tightness is v_t / ES_t.

Job openings for the national-level analysis come from the Job Openings and Labor Turnover Survey (JOLTS). JOLTS data are available beginning in 2001Q1; we use the series constructed by Davis, Faberman and Haltiwanger (2012) to back cast the job openings series to 1994Q1. Unemployment and effective searcher rates are calculated directly from Current Population Survey (CPS) microdata. The effective searcher calculation uses flows across labor market states and movements from job to job to calculate the job finding rates that are the relative search intensity weights for different population groups. Separate job finding rates are calculated for each of twenty-two groups—thirteen groups among the unemployed; three among those who are out of the labor force but say they want a job; four among others who are out of the labor force; and two among the employed.[6] Following Fujita, Moscarini, and Postel-Vinay (2024), we adjust for a change in the procedures used to collect the CPS data that otherwise would lead to an understatement in job-changing rates among employed individuals in later years.

In our empirical implementation, we assume time-invariant relative job-search intensities based on 2006 job finding rates averaged over the months of that year. We construct these intensities as the ratio of job-finding rates for each effective-searcher group relative to a reference group—workers who were recently laid off and expect to be recalled—whose search intensity is normalized to one. The reference group consists of. Using time-invariant relative search intensities is a conservative approach. It attributes variation in effective searchers entirely

[5] As discussed extensively in AHR2020, our measure of effective searchers is similar to that of the Richmond Fed Non-Employment Index (Hornstein, Kudlyak, and Lange 2014).
[6] See Appendix Table A.1 for a full list of the 22 groups along with their 2006 relative job search intensities.

to changes in the population shares of the 22 searcher groups, rather than also incorporating changes in search behavior within groups, so that our measure of effective searchers increases (falls) when the composition of the population shifts toward groups with higher (lower) relative job-search intensities.[7]

Although absolute job-finding rates are highly variable and strongly procyclical, relative job-finding rates are quite stable. For example, while absolute job-finding rates declined substantially between 2006 and 2010, relative job-finding rates had a correlation of 0.98 across the two years. More broadly, 95 percent of the variation in relative job-finding rates over our sample period reflects differences across the 22 groups, while only 5 percent reflects time-series variation within groups. [8]

Matching efficiency can be inferred given information on the flow of hires for given v_t and u_t. The measure of hires for our baseline analysis is seasonally adjusted monthly data from the JOLTS averaged to produce a quarterly series.[9] Labor market tightness is affected not only by movements along the Beveridge curve and shifts in the Beveridge curve due to changes in matching efficiency but also by shifts related to changes in the flow rates across labor market states and out-of-steady-state dynamics. Measures of the relevant flow rates and unemployment changes are calculated from CPS microdata. For all of the series based on CPS data, we construct

[7] Building on the methodology of Hall and Schulhofer-Wohl (2018), AHR2020 constructed group-specific job-finding rates that adjusted for changes in each group's demographic composition over time. Given that results using fixed group-specific job-finding rates were nearly identical, we have adopted the simpler specification here.
[8] Barlevy et al. (2024) suggest that the surge in quits recorded by JOLTS in 2022 and 2023 may be suggestive of increases in the search intensity of employed workers during those years. Direct evidence on this point would be especially valuable. We note, however, that CPS job-to-job flows—adjusted following the method proposed by Fujita, Moscarini and Postel-Vinay (2024) and used as the basis for our measure of relative job-finding rates among the employed—increased much less than quits as measured in the JOLTS. JOLTS quits and CPS job-to-job flows differ not only in the source of the underlying data but also in their conceptual definitions. Further research is needed to reconcile these measures and better understand the discrepancy.
[9] As suggested by equations (11) and (12), information on the unemployment to employment (*ue*) flow rate also can be used to infer what has happened to matching efficiency. As discussed below, we use this alternative as a specification check.

monthly estimates that we seasonally adjust using X-12-ARIMA in EViews and then average to produce quarterly numbers.

In addition to the vacancy and unemployment (effective searcher) measures that we require, the use of the tightness measures in our analysis also depends on the elasticity of the matching function. For this purpose, we draw from our estimates in AHR2020 (see Appendix B of that paper for details). To briefly summarize, the estimates are based on an instrumental variable approach proposed by Hall and Schulhofer-Wohl (2018). This method relates the job-finding rates for specific groups to vacancy duration (v_t/h_t) from JOLTS and yields an estimate of elasticity of the job finding rate with respect to vacancy duration that implies an estimated matching function elasticity. [10] We produce separate estimates for the standard model with only one group of searchers (the unemployed) and the generalized model with 22 groups of searchers. The implied value of α is 0.49 for the standard matching function and 0.57 for the generalized matching function These estimates are reasonably similar and well within the middle of the range of estimates in the matching function literature (Petrongolo and Pissarides 2001). Our main results are broadly unchanged if we apply a common value of the matching function elasticity within the range of the separate estimates.

For the national Phillips curve estimates, following Ball, Leigh and Mishra (2022), our dependent variable is the inflation gap, defined as the difference between observed inflation and expected inflation. Our baseline analysis measures inflation using the seasonally-adjusted core PCE price index (excluding food and energy). This index is monthly; we average across months to produce quarterly estimates and then compute the four-quarter inflation rate as the percentage

[10] This method uses the relationship $f_{it} = \gamma_i A_t (T_t)^{\eta}$ where f_{it} is the job finding rate for group and T_t is mean vacancy duration. This is related to our discussion of equation (10).

change in the quarterly index from *t-4* to *t*. We use the Livingston survey forecast of the average rate of CPI inflation over the next 10 years to measure inflation expectations. The Livingston survey is conducted once every six months; we interpolate between observations to form a monthly expectations series and then average the monthly numbers to produce a quarterly series. The inflation gap measure in our baseline analysis is the difference between four-quarter core PCE inflation and the current quarter Livingston expectations measure. In a sensitivity analysis that replicates our baseline analysis, we replace core PCE with the Cleveland Federal Reserve Bank's median CPI inflation measure.

The tightness measures (and their components) included as explanatory variables in the Phillips curve regressions are average values over the four quarters from *t-4* through *t-1*. We apply a z-score transformation to each of the series used in the Phillips curve estimation (subtracting the mean and dividing by the standard deviation). This transformation does not affect the share of the variation in the dependent variable explained by the different specifications but facilitates comparisons among them.

In a supplemental MSA-level analysis, we make use of data for the 17 MSA's for which we have the data required to construct Beveridge curves, labor market tightness measures and measures of inflation.[11] Our vacancy proxies for these analyses are the MSA-level Help Wanted Index (HWI) series constructed by Barnichon and Shapiro (2024). Prior to 2005, the Conference Board HWI was based on newspaper help wanted ads; a HWI based on online job advertisements was introduced in May 2005 and ran in parallel with the newspaper ad index through July of 2008. To create their series, Barnichon and Shapiro spliced the two sets of HWI indexes together using the method described by Barnichon (2010). Barnichon and Shapiro kindly

[11] These MSAs are Baltimore, Boston, Chicago, Dallas, Denver, Detroit, Houston, Los Angeles, Miami, Minneapolis, New York City, Philadelphia, San Diego, San Francisco, Seattle, St Louis and Washington, DC.

shared their MSA-level HWI with us. We use information from 1995Q4 through 2022Q3, the latest quarter in their dataset.[12]

The MSA-level unemployment rates and effective searcher measures we need for this analysis are constructed using CPS micro data. Because the MSA samples in the CPS are much smaller than the national sample, the MSA-level effective searcher measure is based on four groups—short term unemployed, long term unemployed combined with out of the labor force but want a job, out of the labor force and do not want a job, and currently employed—rather than the 22 groups used for the national-level measure. The MSA-level job finding rates for the four groups used to construct the MSA-level measure are averages for 2004-2006. We use the seasonally-adjusted core four-quarter CPI (inflation rate from *t-4* to *t*), also kindly provided by Barnichon and Shapiro, minus Livingston inflation expectations as our MSA-level inflation gap measure. Following Barnichon and Shapiro (2024), our MSA-level Phillips curve equations include a variety of controls to capture MSA specific effects and time trends.

IV. Beveridge Curve Decompositions

In this section, we compare the standard and generalized measures of labor market tightness along several dimensions. First, we illustrate the Beveridge curves implied by each approach. Next, we examine how matching efficiency differs under the standard and generalized frameworks. Finally, using the decomposition approach developed in Section II, we analyze the sources of shifts in labor market tightness under each approach.

Beveridge Curve and Tightness: Standard vs. Generalized

[12] We intend to update these series and extend our MSA-level analysis through 2025Q2 but have not yet done so.

In the standard framework, labor market tightness is represented as the ratio of vacancies to unemployment. Changes in labor market tightness may occur either because of movements along the standard Beveridge curve (the downward-sloping relationship between u_t and v_t) or because the curve itself shifts inwards or outwards. The standard framework is mis-specified, however, both because the number of unemployed people is an imperfect proxy for the availability of workers to fill an employer's jobs and because a simple count of job openings is an imperfect proxy for the effective level of vacancies. Addressing the former issue, our empirical analysis builds on a measure of effective searchers that accounts not only for changes in the size and composition of the pool of unemployed people but also for fluctuations in the size of the various subgroups among those out of the labor force or currently employed from which an employer might hire. One of our goals is to contrast the generalized Beveridge curve that captures the relationship between job openings and our measure of effective searchers with the standard Beveridge curve that captures the relationship between job openings and unemployment.

The top panel of Figure 1 shows the standard Beveridge curve and the bottom panel shows the generalized Beveridge curve, both plotting data over the period from 1994:Q1 through 2025:Q2. In the top panel, the horizontal axis measures the log of the ratio of unemployment to population age 16 plus; the horizontal axis in the bottom panel measures the log of the ratio of effective searchers to the same population. The vertical axis in both figures is the log of vacancies divided by the population age 16 plus.

Both the standard and the generalized Beveridge curves have a negative slope, but the generalized curve is much steeper than the standard curve. Because the unemployed constitute only a subset of effective searchers, all else equal, an increase in unemployment leads to a

proportionally smaller increase in the pool of individuals actively competing for jobs. Put differently, many groups within the working-age population are potential hires, so a rise in unemployment does not expand the set of potential matches as much as the standard model assumes.

Both curves also exhibit cyclical loops, with vacancies rising relative to both unemployment (in the top panel) and effective searchers (in the bottom panel) during the recovery from the 2007-2009 recession and especially during the recovery from the pandemic recession. During these periods, vacancies rose by more relative to unemployment (in the top panel) and effective searchers (in the bottom panel) than would have been the case along a stable Beveridge curve.

Figure 2 plots the changes in labor market tightness implied by the movements along and shifts in the standard and generalized Beveridge curves in Figure 1. Labor market tightness is shown as log deviations from the mean value of the series over our estimation period. The standard and generalized series generally agree as to whether the labor market is weakening or tightening, moving up and down together, and both reached series peaks following the pandemic. Despite these similarities, reflecting the greater volatility of unemployment as compared to effective searchers, the standard tightness measure is substantially more volatile than the generalized measure. This is especially noticeable during the 2007-2009 downturn, when the standard tightness measure fell more sharply than the generalized measure, as well as during the subsequent recovery, when the standard measure rose more sharply. Though the qualitative behavior of the two series during the pandemic downturn and subsequent recovery followed the pattern of the previous recession and recovery, the gap that opened between the two tightness measures during the pandemic episode is never as large as the gap at the trough of the 2007-2009

recession. Moreover, in the early recovery from the pandemic the generalized measure did not lag behind the standard measure as it did in the recovery from the 2007-2009 recession.

Following the literature (e.g., Ball, Leigh and Mishra 2022; Barnichon and Shapiro 2024), the analysis of the relationship between inflation and tightness in Section V uses four-quarter moving averages of the tightness measures on the right-hand side of the Phillips curve equation. We also use four-quarter moving averages when looking at the effects of the different components of generalized tightness on inflation. To facilitate the connection between the inflation analysis and the decomposition analysis in the next few subsections, we shift for what follows to analyzing four-quarter moving averages. As shown in Figure A.1, the patterns in the moving averages for the two tightness measures are very similar to the patterns depicted in Figure 2.

Decomposition of Standard Matching Efficiency

Holding other factors constant, changes in matching efficiency shift the Beveridge curve and thus the level of vacancies relative to unemployment (in the standard model) or effective searchers (in the generalized model). A number of recent papers working within the standard framework have addressed matching efficiency and its effect on hires and job finding rates (see, e.g., Barlevy et al. 2024, Davis, Faberman and Haltiwanger 2013, and Hall and Schulhofer-Wohl 2018) As shown in equation (9), one of the factors contributing to the variation in standard matching efficiency is variation in the ratio of effective searchers to unemployment. As that ratio increases (decreases), holding constant the other terms in the equation, standard match efficiency rises (falls).

The results of applying the exact decomposition in equation (9), reported in the form of log deviations from overall means, are shown in Figure 3. Although standard and generalized matching efficiency behaved fairly similarly from 1994 through the onset of the Great Financial Crisis, that has not been the case in the ensuing years. Standard matching efficiency fell by more than generalized matching efficiency during the 2007-2009 downturn, reflecting a decline in the ratio of effective searchers to unemployment that added to the effects of the smaller decline in generalized matching efficiency. In AHR2020, we showed that this decline in the ratio of effective searchers to unemployment played a major role in explaining why job filling rates did not rise as much—and job finding rates did not fall as much—as the standard model implied they should have done. Since that time, the ratio of effective searchers to unemployment has continued to fluctuate, rising during the recovery from the Great Financial Crisis, falling during the pandemic downturn and then rising again during the pandemic recovery before declining slightly over the past few years. [13]

As already noted, the fact that the unemployed are only a subset of effective searchers means that effective searchers generally are less cyclical than unemployment. In addition, in the episodes we study, the composition of effective searchers changed over the cycle. During the Great Financial Crisis and its aftermath, the share of long-term unemployed increased substantially and remained elevated. Because long-term unemployed workers search less intensively than the short-term unemployed, effective searchers recovered more slowly than the

[13]In both the 2007-2009 and pandemic recessions and recoveries, the third term in the decomposition, the product of the difference in the estimated elasticity between the standard and the generalized matching function elasticity times the log of the standard labor market tightness measure, also makes a modest but less important contribution to explaining the difference between standard and generalized matching efficiency. As shown in Figure A.2, the decomposition results are very similar if we assume an elasticity of 0.57 for both the standard and the generalized matching function.

unemployment rate itself would suggest, leading the standard model to misinterpret these dynamics.

A somewhat different compositional mechanism operated during the pandemic recession and rapid recovery. The large number of short-term unemployed—especially short-term unemployed on temporary layoff—meant that the ratio of effective searchers to unemployment was unusually high in the early part of the recovery.[14] The standard model read this as an improvement in matching efficiency, when in fact generalized matching efficiency declined.

Decompositions of Standard Labor Market Tightness: Alternative Approaches

The standard matching efficiency decomposition analysis suggests that the standard matching efficiency measure is mis-specified – instead of capturing matching efficiency, much of the variation in the standard measure reflects changes in the ratio of effective searchers to the unemployed. In this section, we use the decompositions of standard labor market tightness from section II to shed further light on the relevance of this mis-specification. In Figure 4, using the decompositions in equations (14) and (15), we compare the actual movements in log standard labor market tightness to the first order log approximations to those series. The approximate decomposition for standard tightness based on equation (14) using the standard model significantly understates the sharp decline in tightness during the Great Financial Crisis. It also deviates noticeably from actual standard tightness during the recovery from the Great Financial Crisis and during the pandemic recession and recovery. In contrast, the approximate

[14] The share of unemployment accounted for by temporary layoffs rose from 16% in 2019Q4 to 73% in 2020Q2 and 43% in 2020Q3 but fell to 26% in 2020Q4 and to 13% in 2021Q1.

decomposition for standard tightness based on equation (15) using the generalized model tracks its target measure relatively closely.[15]

The relatively poor performance of equation (14) for approximating the behavior of standard labor market tightness reflects the limitations of using the standard model to learn about matching efficiency. We have inferred movements in standard matching efficiency from the relationship of JOLTS hires to JOLTS job openings and unemployment, as shown in equation (1). Similarly, we inferred generalized matching efficiency from the relationship of JOLTS hires to JOLTS job openings and effective searchers, as shown in equation (4). An alternative would have been to start with equations (11) and (12), inferring standard matching efficiency from the relationship of *ue* flows to the variables on the right-hand side of equation (12) and generalized matching efficiency from the relationship of *ue* flows to the variables on the right-hand side of equation (11). If the standard model were correct, these two approaches should be capturing the same thing. As can be seen in the left-hand panel of Figure 5, however, the two alternative estimates of matching efficiency for the standard model behave quite differently. In contrast, as shown in the right-hand panel of Figure 5, the two measures of generalized matching efficiency behave very similarly. We view these results as suggesting that the standard model is mis-specified, whereas the generalized model holds up well. In addition, these results suggest that in decomposing standard labor market tightness (v_t/u_t) one should use the generalized model.[16]

[15] The RMSE for the equation (14) approximation for standard tightness is 0.202; the RMSE for the equation (15) approximation is 0.092, less than half as large. In Appendix Figure A.3, we show the components of the decomposition from equation (14) and (15). In the left panel, the terms evidently do not approximate the total sum, whereas in the right panel, they do.

[16] Whereas the decomposition based on equation (1) does a relatively poor job of tracking labor market tightness, as shown in Appendix Figure A.4, the decomposition based on equation (11) closely approximates actual standard labor market tightness.

Decompositions of Generalized Labor Market Tightness

Turning to the decomposition of generalized tightness, Figure 6 compares actual generalized labor market tightness with the approximation given in equation (16). Except for the sharp disruptions during the height of the pandemic in 2020, the generalized tightness approximation tracks the actual data closely.[17] This strong correspondence supports using the decomposition to assess the contributions of the right-hand-side terms of equation (16) to movements in generalized tightness.[18]

As shown in Figure 7, three components drive fluctuations v_t/ES_t : 1) movements along the generalized Beveridge curve, 2) changes in generalized matching efficiency, and 3) other factors that shift the generalized Beveridge curve. Although the decomposition is not exact, outside of a short time during the pandemic period, the sum of these components closely matches observed movements in generalized tightness. Much of the variation reflects movements along the generalized Beveridge curve, but changes in matching efficiency and other structural shifts also play meaningful roles. During the Great Financial Crisis and its aftermath, matching efficiency declined, shifting the generalized Beveridge curve outward. Matching efficiency fell again during the rapid post-pandemic recovery, implying a similar outward shift.

The right-hand panel of Appendix Figure A.5 examines the sources of changes in effective searchers underlying these dynamics. Variations in short-term and long-term unemployment account for about 54 percent and 24 percent of the movement in effective searchers, respectively, while roughly 10 percent reflects changes in the number of individuals

[17] The correlation is 0.97 and the RMSE is 0.085.

[18] Appendix Figure A.4 shows that the decomposition of standard labor market tightness using the generalized model with matching efficiency measured from equation (12) also closely tracks actual labor market tightness. Indeed, that version of the decomposition performs slightly better than the version using generalized matching efficiency from (4), suggesting that matching efficiency could be different across groups. This is an interesting idea that we leave for future work.

out of the labor force who report wanting a job. Additional shifts in the generalized Beveridge curve—separate from generalized matching efficiency—also have contributed to fluctuations in generalized tightness, and at times the effects of generalized matching efficiency have offset the influence of the other factors. For example, in the late 1990s, movements along the generalized Beveridge curve and shifters other than changes in generalized matching efficiency increased tightness, while improved generalized matching efficiency pushed in the opposite direction. After the Great Financial Crisis, declines along the curve and other shifters reduced tightness, but weaker generalized matching efficiency raised the vacancy-to-effective-searcher ratio relative to what it otherwise would have been. Generalized matching efficiency has likewise remained low since the pandemic, exerting upward pressure on generalized tightness.

Taking stock, we draw three main conclusions. First, standard labor market tightness is substantially more volatile than generalized labor market tightness. Second, fluctuations in standard matching efficiency often reflect changes in the ratio of effective searchers to the unemployed rather than "true" shifts in matching efficiency, warranting caution in interpreting the standard matching efficiency measure. Third, generalized matching efficiency follows a pattern that is distinct from that of standard matching efficiency, declining much less during the Great Financial Crisis and its aftermath and falling—rather than rising—in the post-pandemic period.

V. **The Relationship Between Labor Market Tightness and Inflation**.

We now turn to the implications of these findings for the relationship between labor market tightness and inflation, examining how changes in alternative measures of labor market tightness and the components that drive those changes relate to inflation dynamics. Several

recent papers have found that the ratio of vacancies to unemployment—what we refer to as the standard tightness measure—outperforms the unemployment rate and other indicators of labor market tightness in a Phillips curve equation (e.g., Furman and Powell 2021; Ball, Leigh and Mishra 2022; Barnichon and Shapiro 2024). We have argued that generalized labor market tightness—vacancies divided by effective searchers—is a conceptually more appropriate indicator of labor market conditions. To the extent that this is so, we would expect generalized tightness to outperform the standard v_t/u_t measure in the Phillips curve equation. Using our generalized tightness measure, Barnichon and Shapiro (2024) provide evidence that supports this view. We extend their analysis by considering nonlinearities in the effects of generalized tightness on inflation and by exploring whether shifts in generalized tightness that occur for different reasons have different implications for inflation.

Before proceeding, we should acknowledge several limitations of the analysis in this section. First, the sample period for which we have the data needed to carry out the analyses we report using national level data is relatively short (1994Q1 through 2025Q2) and contains only one episode of sustained high inflation. In addition, there are many issues raised in the voluminous literature on the Phillips curve that we do not address here. Rather than reporting results for the plethora of available measures of inflation and expected inflation, we focus on commonly used measures (core PCE and median CPI for inflation, the Livingston 10-year index for expected inflation). We explore simple OLS specifications of the relationship between the inflation gap and measures of labor market tightness. We view our results as making a suggestive case —on both theoretical and empirical grounds—for incorporating the generalized measure of tightness into analyses of inflation dynamics and also for rethinking the idea that it is only the level of overall labor market tightness that matters for these dynamics.

As emphasized in section IV, there are multiple reasons why the generalized tightness measure might change. It could move because of movements along the generalized Beveridge curve; because of shifts in the generalized Beveridge curve due to changes in generalized matching efficiency; or because of shifts in the generalized Beveridge curve for other reasons. In addition to comparing the effects of generalized and standard tightness on inflation, using this decomposition allows us to compare the effects on inflation of changes in generalized tightness that arise for different reasons. This analysis helps to clarify the channels through which labor market conditions influence inflationary pressures.

Our baseline linear empirical Phillips curve equation is given by:

(17) $\pi_t - \pi_t^e = \alpha + \beta\theta_t^i + \varepsilon_t$

where the term on the left-hand side of the equation is the inflation gap between actual and expected inflation with $i = s$ for the standard tightness measure and $i = g$ for generalized tightness. Our baseline inflation measure is the core PCE price index (excluding food and energy), measured as a four-quarter change (*t* minus *t-4*). Like Ball, Leigh and Mishra (2022), rather than including expected inflation on the right-hand side of the equation, we analyze the relationship between the inflation gap (the difference between actual inflation and expected inflation measured using the Livingston 10-year index) and measures of labor market tightness. This is equivalent to imposing a coefficient of 1.0 on expected inflation on the right-hand side of the equation. The tightness measures on the right-hand side of the equation are average values over the four preceding quarters (*t-4* through *t-1*). As a reminder, to facilitate comparisons across models, all of the tightness measures used in estimation are z-score transformed.

The first column of Table 1 reports the results of estimating a linear model that relates the inflation gap to standard labor market tightness; the third column reports the results of estimating

a linear model with generalized tightness on the right-hand side. Figure 8 shows the actual inflation gap alongside fitted values from the column (1) and (3) regressions. In both equations, the coefficient on the tightness measure is positive and statistically significant, but the coefficient in the regression that includes the generalized tightness measure is larger and that model explains more of the variance in the inflation gap. Although the fitted values from the two regressions are similar over much of our time period and neither fitted series perfectly captures the movements in inflation, the generalized measure does a better job of explaining the post-pandemic surge in inflation than the standard measure.

To allow for potential nonlinearities in the relationship between labor market tightness and inflation, we estimate specifications in which tightness enters through a natural cubic spline rather than linearly. The spline specification replaces the single tightness regressor with a set of piecewise-cubic polynomial basis functions joined smoothly at pre-specified knot values, allowing the marginal effect of tightness on inflation to vary flexibly across the distribution of tightness while maintaining continuity and smooth first and second derivatives. This approach provides a flexible semi-parametric approximation to an unknown functional form without imposing the strong curvature restrictions inherent in quadratic or higher-order polynomials (de Boor 1978; Stone 1985; Hastie, Tibshirani, and Friedman 2009; Harrell 2001).

Model selection criteria strongly favor the spline specification— Wald χ^2 tests reject the linear model in favor of the spline, and both AIC and BIC are lower for the spline models, indicating improved fit despite the additional parameters. The latter is also evident in the adjusted R-squareds. Because spline coefficients correspond to basis functions rather than economically meaningful slopes, the individual parameter estimates are not directly interpretable; we report them for completeness, but evaluate the relationship using fitted-value

plots and associated confidence intervals, which transparently display how the predicted inflation gap varies with labor market tightness over its observed range. Figure 9 illustrates the nonlinear relationships between the inflation gap and tightness implied by the estimated regression coefficients. Figure 10 plots the actual inflation gap alongside fitted values from these regressions.[19]

The cubic spline models imply a relationship between labor market tightness and inflation that is highly nonlinear, such that the predicted inflation gap responds more strongly to increases in tightness when tightness is already high. This can be seen clearly in Figure 9. The nonlinearity implied by the generalized tightness model is more pronounced and, as indicated by the width of the error bands, more precisely estimated than that in the model based on standard tightness.

Comparing the cubic spline specifications, the adjusted R-squared for the nonlinear generalized tightness model is substantially higher, and the AIC and BIC substantially lower, than for the nonlinear standard tightness model. As was true for the linear specifications, the fitted values based on the model that uses generalized labor market tightness track the movements in inflation movements more closely than those for the version that uses standard tightness. This is particularly true during the period of elevated inflation following the pandemic, when inflation rose sharply alongside historically high vacancy rates but comparatively muted movements in effective searchers.

The finding of substantial nonlinearity in the relationship between inflation and labor market tightness is consistent with other recent work that has documented nonlinearities in that relationship (Ball, Leigh and Mishra 2022; Benigno and Eggertsson 2024). What is new about

[19] As a robustness check, we have explored cubic splines with 3 knots rather than 4 knots. The AIC and BIC criteria support the use of 4 knots.

our results is showing that these nonlinearities are even stronger when the model captures conditions in the labor market using a generalized tightness measure (vacancies divided by effective searchers) rather than the standard tightness measure (vacancies divided by unemployment).

To better understand how and why generalized tightness drives inflation, we next estimate models that allow different sources of movements in generalized tightness to affect the inflation gap differently. As previously noted, changes in generalized tightness can be decomposed into three components—movements along the generalized Beveridge curve, shifts attributable to changes in generalized matching efficiency, and other shifts. The model in column (1) of Table 2 reproduces the linear generalized tightness specification from Table 1, while the model in column (2) replaces the aggregate measure with these three components entered separately (as defined in equation 16). Allowing the components to enter individually increases the model's explanatory power, indicating that the sources of tightness and not simply the overall level of tightness matter for inflation.[20]

The estimated coefficients suggest that variation in generalized matching efficiency is the most important contributor to inflation dynamics, followed by other Beveridge-curve shifts, with movements along the curve (changes in effective searchers) playing a comparatively smaller role. This pattern is broadly consistent with Barnichon and Shapiro (2024), who find that changes in standard matching efficiency are particularly relevant for inflation. As discussed in Section IV, however, the decomposition based on the standard Beveridge curve conflates "true"

[20] Because the decomposition in equation (16) is not exact, the three components reported in Table 2 do not fully account for the variation in generalized tightness. As a robustness check (results not reported), we re-estimated the specification in the right column of Table 2 including the residual component of generalized tightness as an additional regressor. The coefficient on this residual term is economically small and statistically insignificant, and the estimated coefficients on the three reported components remain virtually unchanged.

(generalized) changes in matching efficiency with changes in the ratio of effective searchers to unemployment. The generalized Beveridge-curve framework used here separates these mechanisms, allowing us to isolate variation in generalized matching efficiency that is not driven by changes in the composition of job seekers.

While the coefficient estimates are informative about marginal effects—i.e., how inflation responds to a change in one component holding the others fixed—in practice movements along and shifts in the Beveridge curve tend to occur jointly. To provide a complementary perspective that reflects this empirical co-movement, we report a Shapley–Owen decomposition of the model's R^2 (see, e.g., Huettner and Sunder 2012; Mueller and Spinnewijn 2025). This procedure allocates the regression's explained variation across the three components based on their incremental contribution to model fit, averaged over all possible orderings in which the variables could be introduced.

The Shapley-Owen decomposition indicates that changes in generalized matching efficiency account for roughly three-fifths of the model's explained variation in inflation, with smaller but nontrivial contributions from Beveridge-curve shifts and movements in effective searchers. These results reinforce the message of the coefficient estimates—inflation responds most strongly to forces that alter the efficiency with which vacancies are translated into hires, and less to fluctuations in the effective supply of job seekers associated with cyclical movements along the Beveridge curve.

Further pursuing the question of the relationship between generalized labor market tightness and inflation, Figure 11 compares predictions from the cubic spline model using generalized tightness with predictions based on the decomposition just described. The adjusted R-squared in the decomposed model (0.533) is not as large as that for the cubic spline in

generalized tightness (0.663), but as can be seen in Figure 11, the decomposition model captures a good deal of the inflation dynamics captured by the nonlinear model.

As a final exercise, we extend the analysis by combining the nonlinear spline specifications and decomposition approaches, fitting a separate cubic spline for each of the three decomposition terms. The estimated coefficients from this model are not readily interpretable and we have relegated them to Appendix Table A.2. Figure 12 shows the fitted inflation values from this elaborated cubic spine together with the fitted values from the simple cubic spline and the actual inflation values. Summary statistics from the model are reported in Appendix Table A.3. Allowing each component of generalized tightness to enter nonlinearly improves on the nonlinear model's explanatory power, raising the adjusted R-squared from 0.663 for the simple nonlinear model to 0.763 for the elaborated nonlinear model. As in the simpler specification, matching efficiency remains the most important contributor to inflation dynamics and changes in effective searchers associated with movements along the Beveridge curve remain the least important.

Using Median Inflation

Ball, Leigh and Mishra (2022), among others, advocate for using an alternative measure of inflation – specifically median inflation that inherently abstracts from some of the high frequency volatility that is present in the core PCE measure.[21] They also find that there is a stronger empirical relationship between the inflation gap using median inflation and labor market tightness relative to using core PCE inflation. We have replicated all of our analyses using this

[21] Another alternative of interest is the relationship between tightness and wage inflation. The challenge for exploring the latter in this context is the need to control for labor productivity growth as in Ball, Leigh and Mishra (2022). Given the complex dynamics between wages and productivity, we leave this for future work.

alternative inflation measure. Since the results are broadly similar, we relegate the tables and figures to Appendix B and summarize key findings here.

Like Ball, Leigh and Mishra (2022), we find a stronger relationship between the inflation gap and standard labor market tightness when using the median inflation measure than when using the core PCE price index. As with our baseline results, however, the relationship in the median inflation models is even stronger using the generalized tightness measure than when using the standard measure. Focusing on the models using the generalized tightness measure, the adjusted R-squared in the model using the cubic spline specification (0.914) is substantially higher than that for the corresponding linear specification (0.711), indicating that, as for core PCE inflation, the relationship between inflation and tightness using the median inflation measure is nonlinear. Finally, breaking generalized labor market tightness into its component pieces, the coefficient estimates suggest that. holding the values of the other components constant, generalized matching efficiency has the largest effect on inflation.

MSA Level Analysis

Given the relatively small number of quarters of data available for our aggregate analysis, we also have explored the relationship between inflation and labor market tightness using spatial variation across MSAs, following the approach in Barnichon and Shapiro (2024). We build on their work by combining their MSA-level Help Wanted Index (HWI) series with tabulations from the CPS microdata to construct measures of both standard and generalized tightness. Because this exercise pushes the CPS data hard, we cannot implement the 22-group decomposition used in our baseline analysis. Instead, we construct effective searchers using four groups—short-term unemployed, long-term unemployed combined with not in the labor force

but want a job, others not in the labor force and the employed.[22] As can be seen in Appendix Figure A.6, in the national level data, the effective searcher measure using the four-group breakout captures much but not all of the variation in the effective searcher measure using 22 groups. We also cannot decompose tightness into its components because the hires data used for measuring matching efficiency in our baseline analysis are not available at the MSA level.[23]

Given the available data, we estimate simple linear relationships between the MSA-level inflation gap and alternative measures of MSA-level labor market tightness. The dependent variable in these regressions is MSA-level core CPI four-quarter inflation (percent change from quarter *t-4* to *t*) minus expected inflation based on the Livingston data. The key regressor is log labor market tightness, measured at the MSA level, averaged over the prior four quarters (*t–4* through *t–1*) and standardized to facilitate comparison across tightness concepts. Following Barnichon and Shapiro, we include a range of MSA-level controls—most notably MSA fixed effects and MSA-specific time trends—to account for time-invariant local heterogeneity and differential local trends.

A key difference between our analysis and Barnichon and Shapiro's analysis is that, because the CPS data required to construct our generalized tightness measure is not available prior to 1994, we have fewer years of data; their MSA analysis spans 1982–2022, while ours covers 1995-2022. In our setting, including a full set of quarter fixed effects would absorb much of the remaining within-MSA variation in tightness that identifies the relationship of interest. We therefore control for aggregate conditions using a structured set of macro controls—NBER

[22] We also estimated five-group models with people who were out of the labor force but wanted a job separated from the long term unemployed. The results were very similar to the reported four-group results, but some of the cell sizes for these groups were very small.

[23] JOLTS hires data are not available for MSAs and measuring MSA-level hires from the CPS gross flows is pushing the data too hard. Experimental JOLTS data are available at the state level, and we considered an analysis that exploited cross-state rather than cross-MSA variation, but the JOLTS state estimates are modeled estimates and they exhibited relatively little systematic geographic variation.

recession indicators (equal 1 if in Recession, 0 otherwise), a Great Financial Crisis (GFC) dummy (equal 1 if in GFC Recession, 0 otherwise), a pandemic recession dummy (equal 1 if in pandemic recession, 0 otherwise), and national controls including national GDP growth and national four-quarter CPI core inflation, each measured as four-quarter averages over *t–4* through *t–1*, mirroring the timing used for MSA tightness.

Table 3 reports the results. Consistent with our earlier findings, generalized labor market tightness is more strongly associated with inflation than standard tightness. The estimated coefficient on generalized tightness is approximately 25 percent larger than the corresponding coefficient on standard tightness, implying that a one-standard-deviation increase in generalized tightness is associated with a meaningfully larger increase in inflation than an equivalent increase in standard tightness.

Taking Stock of the Evidence on Inflation

Taken as a whole, the results in this section lead to two main conclusions. First, generalized labor market tightness is a more informative predictor of inflation than the conventional vacancy–unemployment ratio. This holds true whether we model the relationship between labor market tightness and inflation as linear or (more consistent with the data) allow for nonlinearities. Second, models in which we allow the different contributors to generalized labor market tightness to affect inflation differently reveal that inflation responds most strongly to changes in generalized tightness due to changes in generalized matching efficiency and other forces that shift the generalized Beveridge curve rather than to fluctuations in effective labor supply.

One possible explanation for this latter finding is that employers view changes in generalized matching efficiency and some of the other factors that shift the generalized Beveridge curve as more likely to persist than changes in the number of effective searchers due to movements along the generalized Beveridge curve. Movements up and down the generalized Beveridge curve are cyclical in nature, meaning that they are unlikely to be permanent. Shifts in the generalized Beveridge curve and perhaps especially shifts due to changes in generalized matching efficiency may be more likely to persist. Consistent with this view, in our data, generalize matching efficiency and the other generalized Beveridge curve shifter term are more persistent (correlations from one quarter to the next of 0.98 and 0.85, respectively) than the effective searcher term (correlation of 0.71).

These findings have direct implications for how analysts and policymakers should interpret Beveridge curve shifts in real time. During periods when generalized tightness is behaving differently from standard tightness, apparent signs of labor market strengthening or weakening based on the standard tightness measure may provide misleading signals with regard to inflationary pressure. Relatedly, variation in measures of matching efficiency derived from the standard model may provide misleading guidance about the inflation risk associated with actual changes in matching efficiency as measured by the generalized model. The generalized framework provides a more nuanced lens through which to assess the inflationary consequences of labor market developments.

VI. **Conclusion**

This paper revisits the measurement of labor market tightness and matching efficiency by enhancing the standard vacancy–unemployment framework. More specifically, we employ a

generalized matching model that accounts for heterogeneity in search intensity and recognizes that hires arise not only from unemployment, but also through flows from nonparticipation into employment and from job-to-job transitions. Doing so yields a set of results that alter how we interpret Beveridge curve shifts and their implications for inflation.

Building on earlier work, we find that generalized labor market tightness—vacancies relative to effective searchers—is substantially less volatile than the conventional vacancy-to-unemployment ratio. Because effective searchers include groups beyond the unemployed, the effective labor supply available to employers fluctuates less sharply over the business cycle than the unemployment rate. As a result, movements in the standard tightness measure often exaggerate cyclical swings relative to the generalized labor market tightness measure.

Second, we show that what is commonly interpreted as variation in "matching efficiency" within the standard model largely reflects movements in the ratio of effective searchers to unemployment rather than changes in "true" matching efficiency (the measure of matching efficiency that emerges from the generalized model). When unemployment rises sharply—as in the Great Financial Crisis—the effective pool of job seekers expands by less than the unemployment rate would suggest. The standard model attributes the resulting divergence between hires and unemployment to a sharp deterioration in matching efficiency. In contrast, the generalized model reveals that generalized matching efficiency declined much less during the Great Financial Crisis and its aftermath. Similarly, in the post-pandemic period, the standard measure of matching efficiency largely recovered from its pandemic-era decline, whereas generalized matching efficiency continued to drift downward. These results suggest that the "matching efficiency" term in the standard framework conflates compositional changes in effective searchers with genuine changes in the efficiency of matching.

We then examine the implications of these measurement differences for inflation dynamics. Across a range of specifications, generalized labor market tightness accounts for more of the variation in the inflation gap than the standard vacancy–unemployment ratio. This is true in a simple linear Phillips curve specification and even more so in nonlinear models that allow the effect of tightness on inflation to increase at high levels of tightness. The nonlinear relationship between tightness and inflation is more pronounced and more precisely estimated when tightness is measured using the generalized concept.

Moreover, when we decompose generalized tightness into movements along the generalized Beveridge curve, shifts due to changes in generalized matching efficiency, and other curve-shifting factors, both the regression coefficients and the Shapley-Owen decompositions of model fit imply that the matching efficiency component is especially important in accounting for inflation dynamics. These results suggest that inflation responds differentially to different sources of variation in generalized labor market tightness, with a stronger response to changes in the efficiency with which vacancies are converted into hires than to cyclical movements in effective labor supply that may be viewed by firms as less persistent.

At the same time, our analysis leaves important avenues for future research. Most notably, while we have generalized the supply side of the matching framework by incorporating heterogeneity in search intensity across worker groups, we have not similarly generalized the demand side. In particular, we have not incorporated measures of effective vacancies that account for variation in employer recruiting intensity, as emphasized in establishment-level analyses. To the extent that recruiting intensity varies over time, some of what we attribute to changes in matching efficiency may instead reflect unmeasured variation in effective vacancies.

In addition, our measure of effective searchers is deliberately conservative. We use time-invariant relative search intensities, so that changes in our measure of effective searchers reflects only compositional shifts across groups with different fixed relative job-finding rates. This approach avoids conflating cyclical movements in job-finding rates with changes in search intensity, but it also abstracts from potential time variation in search effort within groups. Developing and collecting more direct evidence on search intensity across worker groups—and on recruiting intensity across employers—would allow for a richer and potentially more accurate characterization of both sides of the matching process.

More broadly, our findings point to the value of rethinking how labor market conditions are summarized in macroeconomic analysis. The generalized framework developed here provides a more internally consistent account of job-filling, job-finding, and inflation dynamics than the standard vacancy–unemployment approach. By distinguishing "true" changes in matching efficiency from compositional shifts in effective search, it offers a clearer lens through which to interpret Beveridge curve movements and to assess the inflationary implications of labor market developments.

Taken together, the results suggest that incorporating a broader and more behaviorally grounded measure of labor market tightness can sharpen both positive analysis and policy assessment. At the same time, the limitations we have identified underscore that measuring effective search and effective recruiting remains an open and important frontier for research.

References

Abraham, Katharine, John Haltiwanger, and Lea Rendell. 2020. "How Tight is the US Labor Market?," *Brookings Paper on Economic Activity,* Spring, 97-165. https://doi.org/10.1353/eca.2020.0001

Ahn, Hie Joo, and Leland D. Crane. 2020. *"Dynamic Beveridge Curve Accounting,"* Finance and Economics Discussion Series 2020-027. Washington: Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2020.027

Ball, Laurence M., Daniel Leigh, and Prachi Mishra. 2022. "Understanding U.S. Inflation During the COVID Era," *Brookings Papers on Economic Activity*, Fall, 1–80. https://doi.org/10.1353/eca.2022.a901276

Barlevy, Gadi, R. Jason Faberman, Bart Hobijn, and Ayşegül Şahin. 2024. "The Shifting Reasons for Beveridge Curve Shifts," *Journal of Economic Perspectives* 38(2): 83–106. https://doi.org/10.1257/jep.38.2.83

Barnichon, Régis. 2010. "Building a Composite Help-Wanted Index," *Economics Letters* 109(3): 175–178. https://doi.org/10.1016/j.econlet.2010.08.029

Barnichon, Regis and Adam Shapiro. 2024. "Phillips meets Beveridge," *Journal of Monetary Economics*, 148(S), 103660. https://doi.org/10.1016/j.jmoneco.2024.103660

Benigno, Pierpaolo, and Gauti B. Eggertsson. 2024. "*Revisiting the Phillips and Beveridge Curves: Insights from the 2020s Inflation Surge*," National Bureau of Economic Research Working Paper 33095. https://doi.org/10.3386/w33095

Blanchard, Olivier J., and Peter A. Diamond. 1989. "The Aggregate Matching Function," in *Growth/Productivity/Unemployment*, Peter A. Diamond, ed., Cambridge, MA: MIT Press, 159–201.

Blanchard, Olivier J., and Peter A. Diamond. 1992. "The Flow Approach to Labor Markets," *American Economic Review* 82(2): 354–359.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2012. "Labor Market Flows in the Cross Section and Over Time," *Journal of Monetary Economics* 59 (1): 1–18. https://doi.org/10.1016/j.jmoneco.2011.10.001.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2013. "The Establishment-Level Behavior of Vacancies and Hiring," *Quarterly Journal of Economics*, 128(2): 581-622. https://doi.org/10.1093/qje/qjt002

de Boor, Carl. *A Practical Guide to Splines.* Applied Mathematical Sciences, vol. 27. New York: Springer-Verlag, 1978.

Diamond, Peter A. 1982. "Wage Determination and Efficiency in Search Equilibrium," *Review of Economic Studies* 49(2): 217–227. https://doi.org/10.2307/2297271

Fujita, S., G. Moscarini, and F. Postel-Vinay. 2024. "Measuring Employer-to-employer Reallocation," *American Economic Journal: Macroeconomics*, 16(3): 1-51. https://doi.org/10.1257/mac.20210076

Furman, Jason, and William Powell III. 2021. "*What Is the Best Measure of Labor Market Tightness?*" RealTime Economic Issues Watch, Peterson Institute for International Economics, November 22. https://www.piie.com/blogs/realtime-economics/2021/what-best-measure-labor-market-tightness

Hall, Robert E., and Sam Schulhofer-Wohl. 2018. "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Job-Seekers." *American Economic Journal: Macroeconomics* 10(1): 1–32. https://doi.org/10.1257/mac.20170061

Harrell, Frank E. 2001. *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*. New York: Springer-Verlag.

Hastie, Trevor, Robert Tibshirani, and Jerome H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* 2nd ed. New York: Springer, 2009.

Hornstein, Andreas, Marianna Kudlyak, and Fabian Lange. 2014. "Measuring Resource Utilization in the Labor Market." *Federal Reserve Bank of Richmond Economic Quarterly* 100, no. 1: 1–21.

Huettner, Frank, and Marco Sunder. 2012. "Axiomatic Arguments for Decomposing Goodness of Fit According to Shapley and Owen Values," *Electronic Journal of Statistics* 6: 1239–1250. https://doi.org/10.1214/12-EJS710

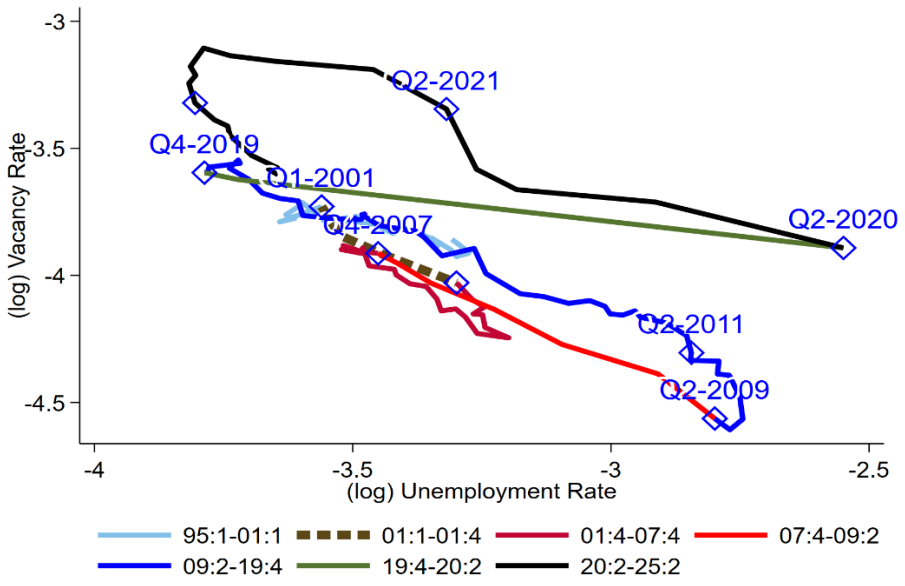
Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies* 61(3): 397–415. https://doi.org/10.2307/2297896

Mueller, Andreas I., and Johannes Spinnewijn. 2025. "The Nature of Long-Term Unemployment: Predictability, Heterogeneity and Selection," *Journal of Political Economy* 133 (12): 3846–3902. https://doi.org/10.1086/738045

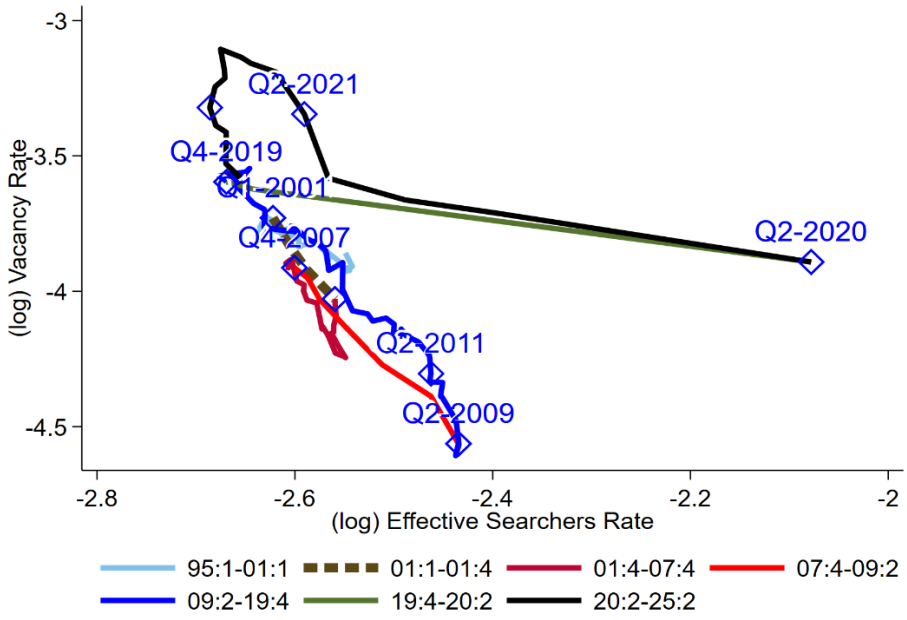
Petrongolo, Barbara, and Christopher A. Pissarides. 2001. "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature* 39(2): 390–431. https://doi.org/10.1257/jel.39.2.390

Pissarides, Christopher A. *Equilibrium Unemployment Theory.* 2nd ed. Cambridge, MA: MIT Press, 2000.

Stone, Charles J. 1985. “Additive Regression and Other Nonparametric Models,” *Annals of Statistics,* 13(2): 689–705. https://doi.org/10.1214/aos/1176349548



(a)



(b)

Figure 1: Standard and Generalized Beveridge Curve

Notes: Authors' calculations using JOLTS and CPS data. All series reported as log of ratio of series (vacancies, unemployment, or effective searchers) to population aged 16 plus. The X-axis scale is different for panels (a) and (b).

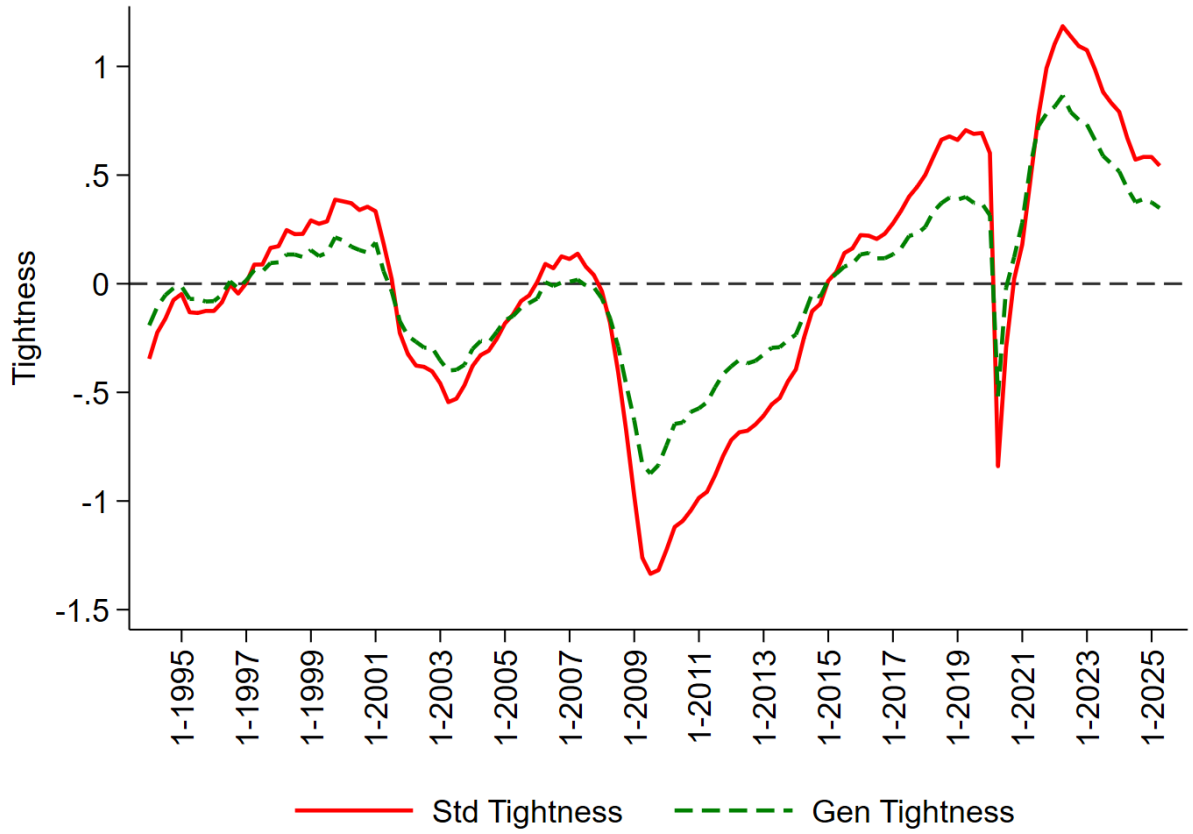
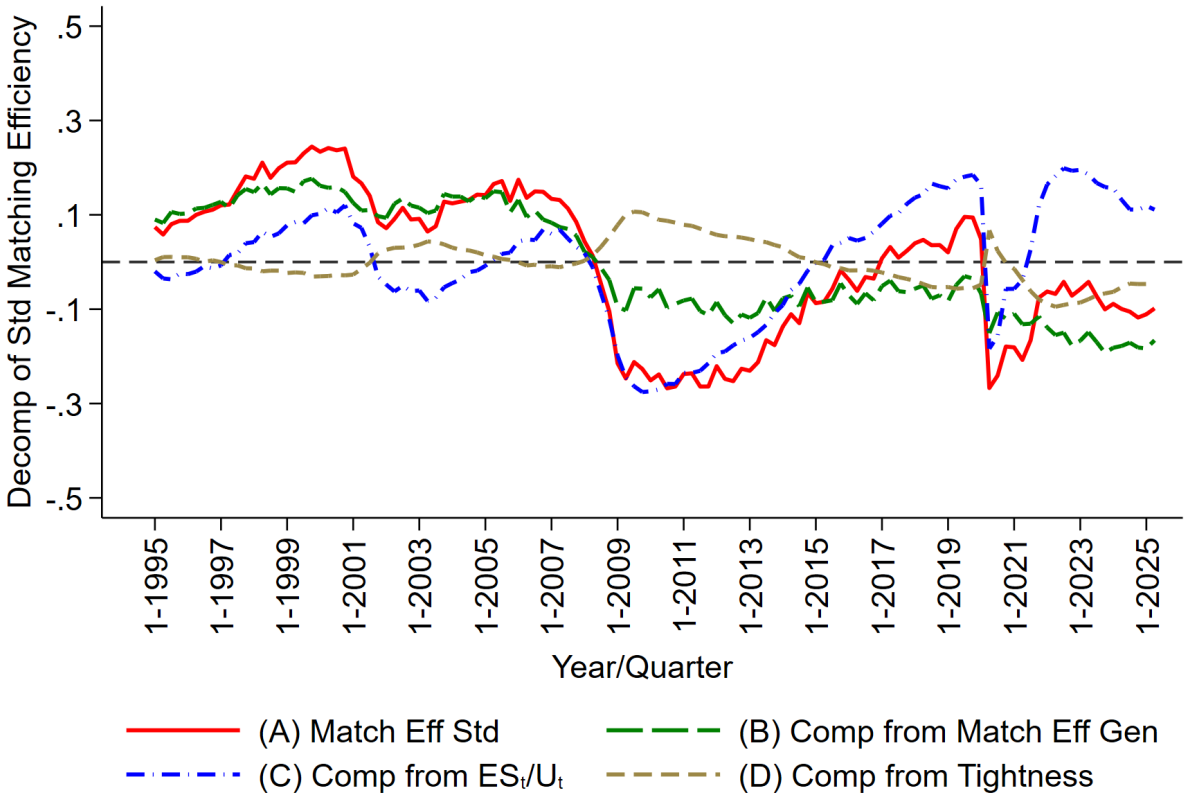


Figure 2: Standard and Generalized Labor Market Tightness

Notes: Authors' calculations using JOLTS and CPS data. All series are demeaned logs. Standard tightness is *v/u* and generalized tightness is *v/ES*.



Note: By construction, Shift (A) = Shift (B) + Shift (C) + Shift (D).

Figure 3: Decomposition of Matching Efficiency

Notes: Authors' calculations using JOLTS and CPS data. Decomposition of log matching efficiency from standard model (Match Eff Std) using equation (9). Terms are a component for generalized matching efficiency (Comp from Match Eff Gen), a component for $\sum_i \rho^{l_i} l_{it} / U_t$ (Comp from ES/U), and a component for different matching elasticities times standard tightness (Comp from Tightness).

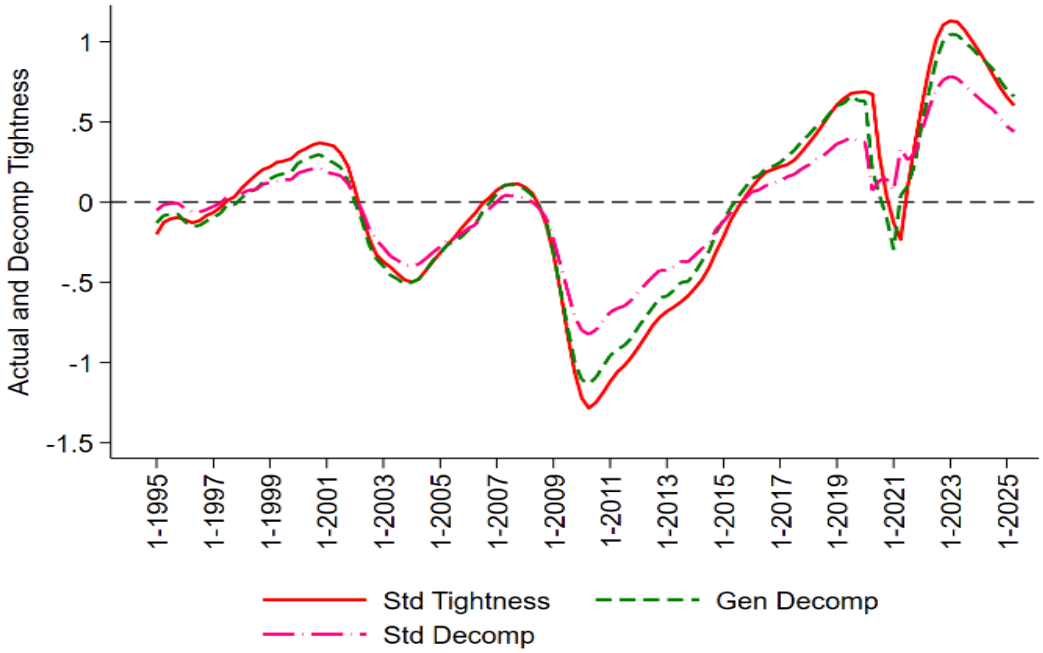


Figure 4: Standard Tightness: Actual and Approximations from Decompositions using Standard Model from (14) and Generalized Model from (15).

Notes: Authors' calculations using JOLTS and CPS data.

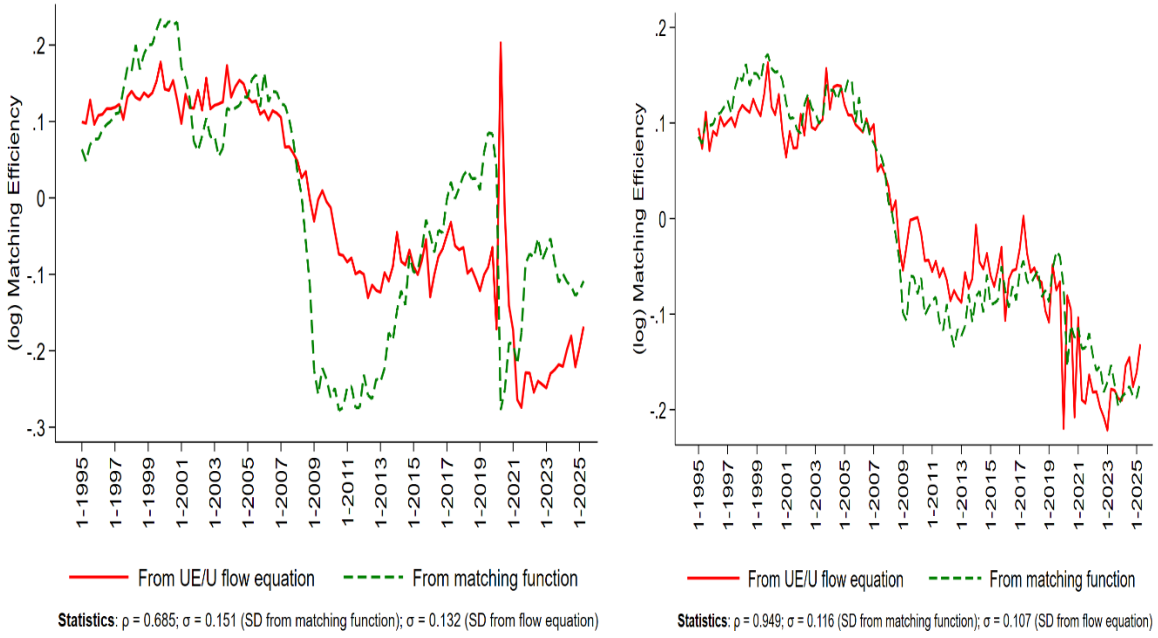


Figure 5: Robustness Check on Matching Efficiency

Notes: Authors' calculations using JOLTS and CPS data. Matching efficiency from UE/U flow equation is from equation (12) for standard model and equation (11) for generalized model. Matching efficiency from matching function is from equation (7) for standard model and equation (8) for generalized model.

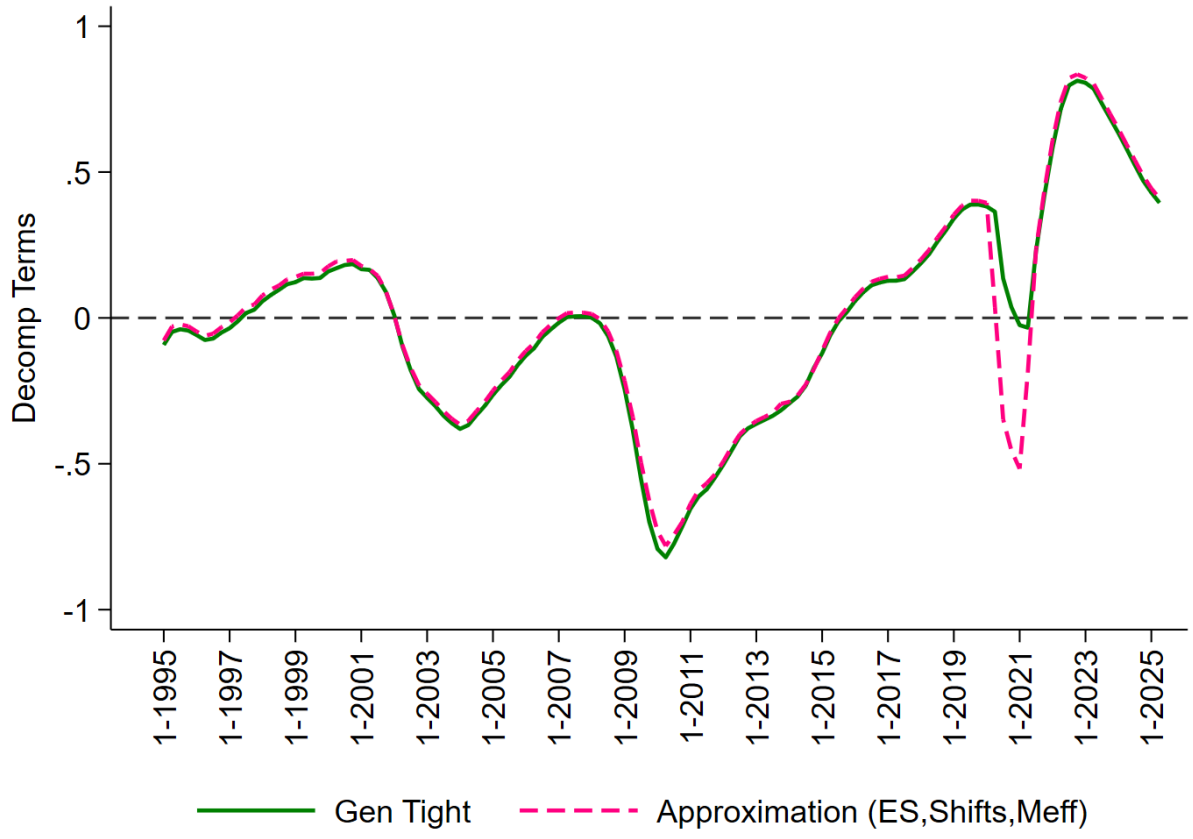


Figure 6: Generalized Labor Market Tightness: Actual and Approximation from Decomposition in equation (16).

Notes: Authors' calculations using JOLTS and CPS data.

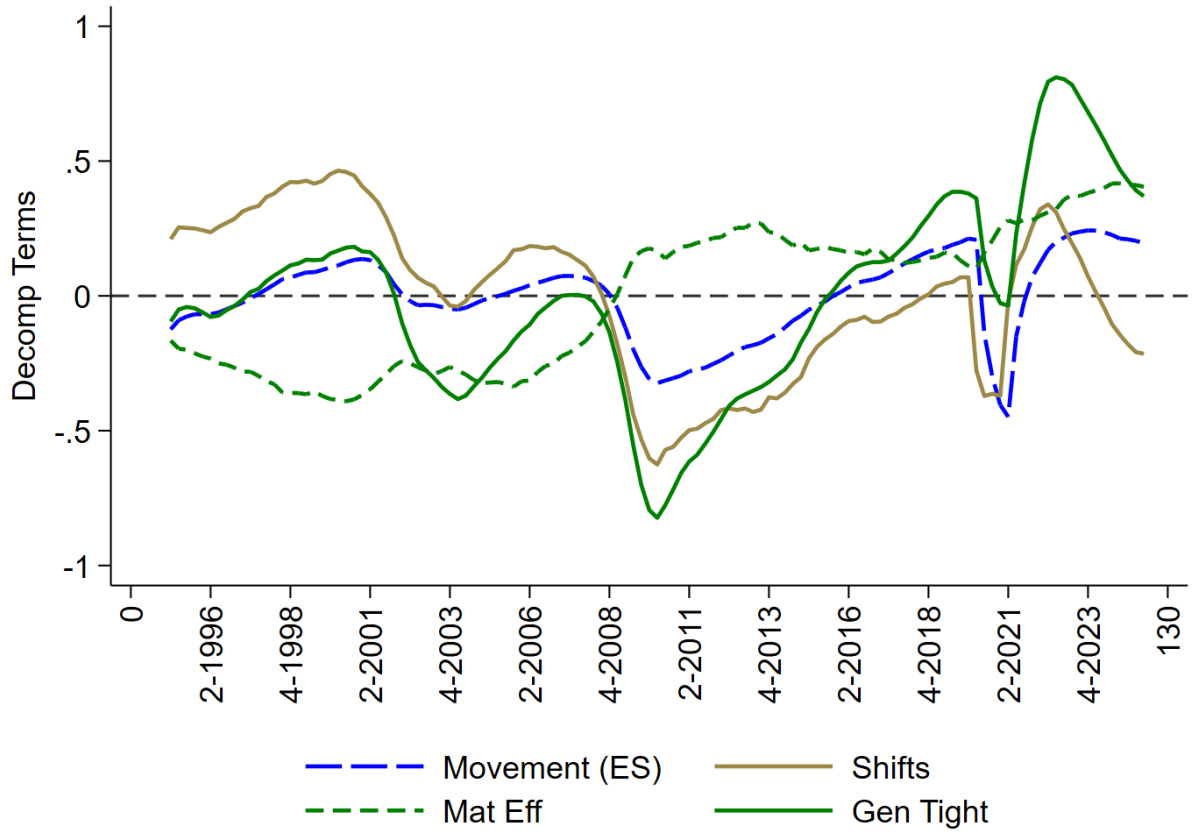


Figure 7: Decomposition of Generalized Tightness (3-way Decomposition from equation (16))

Notes: Authors' calculations using JOLTS and CPS data.

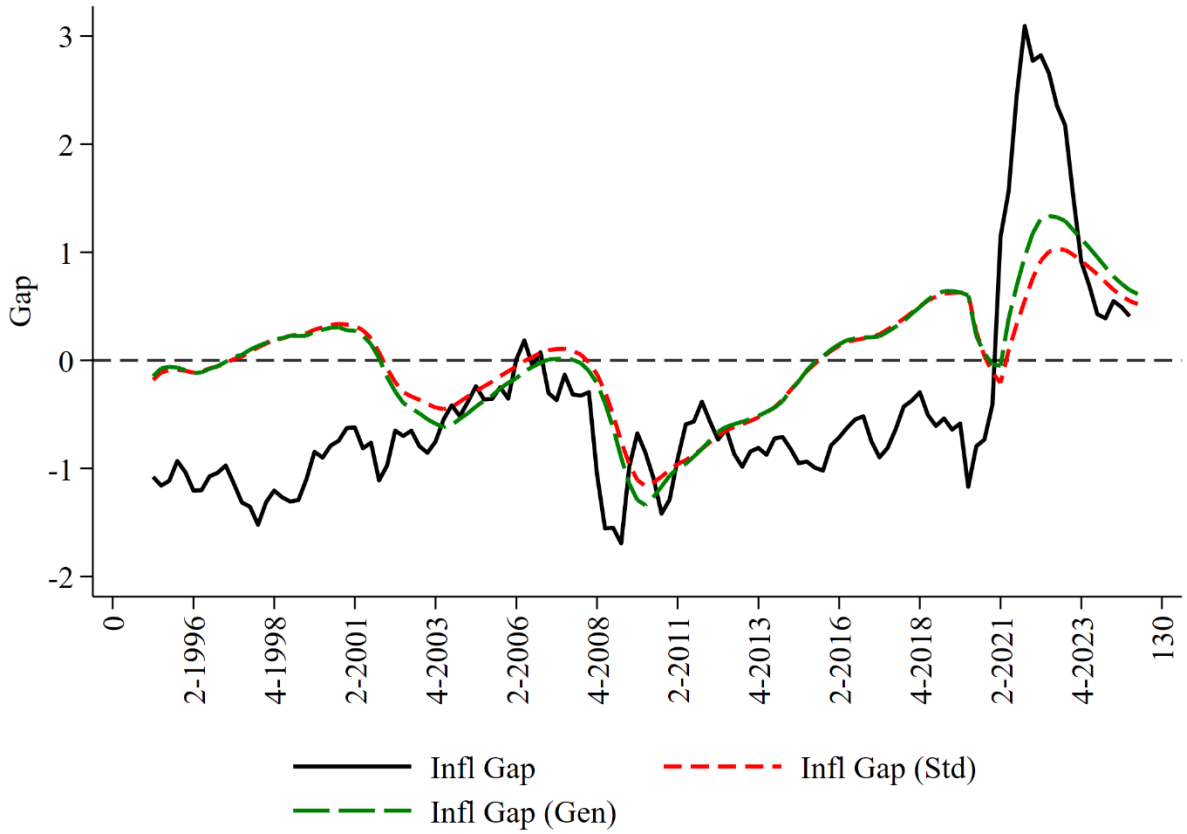


Figure 8: Actual Inflation Gap and Predicted Inflation Gap from Labor Market Tightness (Standard vs Generalized)

Notes: Authors' calculations using JOLTS national data, CPS micro data and the inflation gap measured using PCE core inflation and inflation expectations from the Livingston survey. Predicted inflation gap from column (1) of Table 1 for standard tightness (Infl Gap (Std)) and column (3) of Table 1 for generalized tightness (Infl Gap (Gen)).

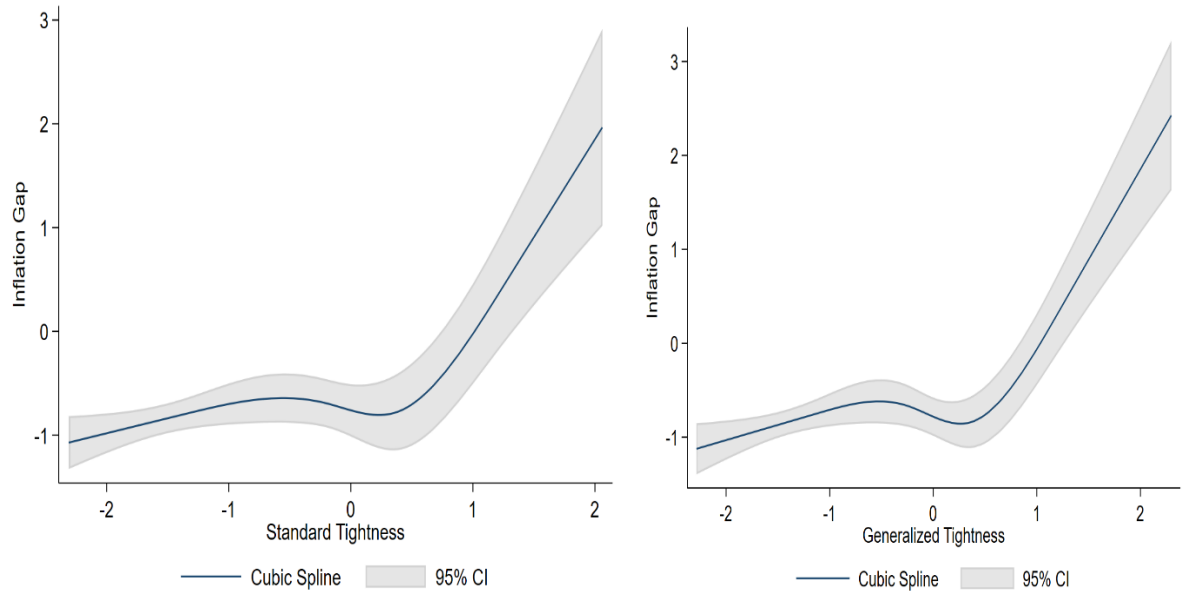


Figure 9: Nonlinear Relationship between Inflation Gap and Labor Market Tightness (Standard vs. Generalized)

Notes: Authors' calculations using JOLTS national data, CPS micro data and the inflation gap measured using PCE core inflation and inflation expectations from the Livingston survey. Figures derived from cubic spline from column (2) of Table 1 for Standard Tightness and column (4) of Table 1 for Generalized Tightness. Confidence intervals produced using the delta method.

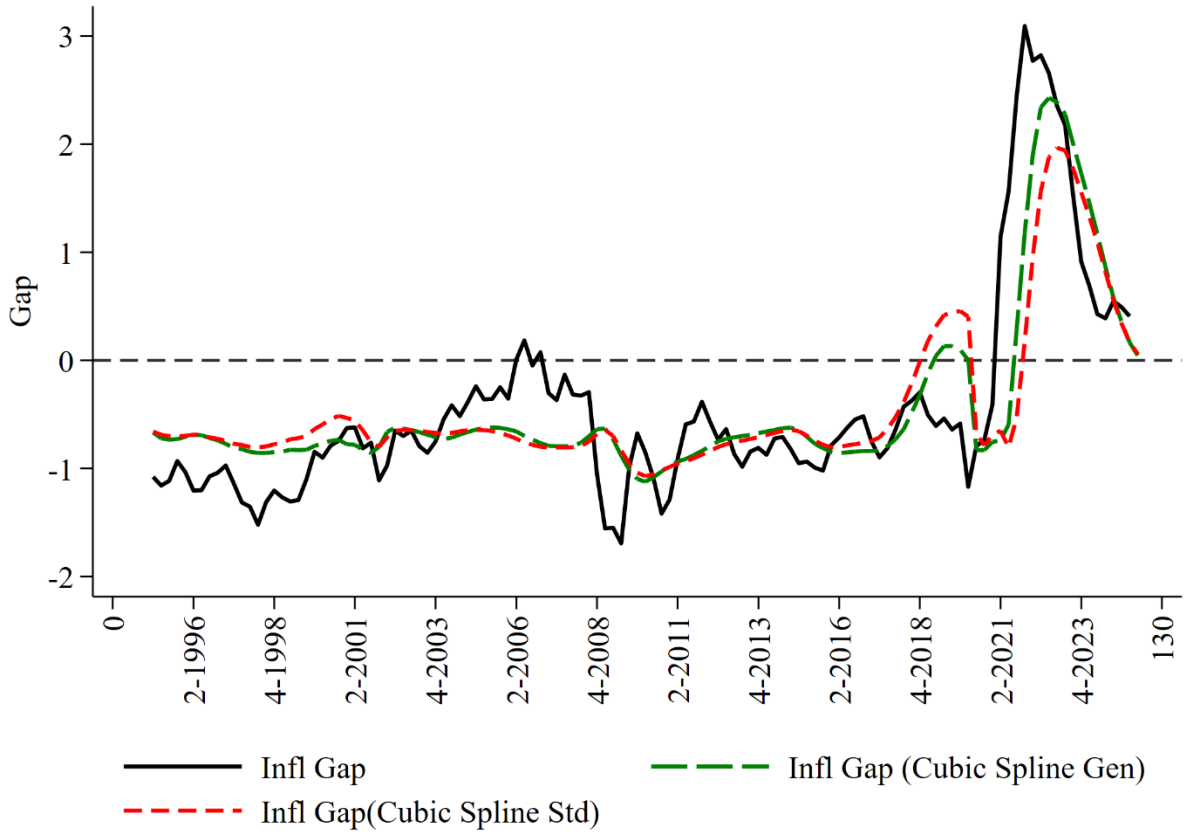


Figure 10: Actual Inflation Gap and Predicted Inflation Gap from Labor Market Tightness (Standard vs Generalized): Cubic Spline Model

Notes: Authors' calculations using JOLTS national data, CPS micro data and the inflation gap measured using PCE core inflation and inflation expectations from the Livingston survey. Predicted inflation gap for Cubic Spline Std from estimates in column (2) of Table 1 and for Cubic Spline Gen from estimates in column (4) of Table 1.

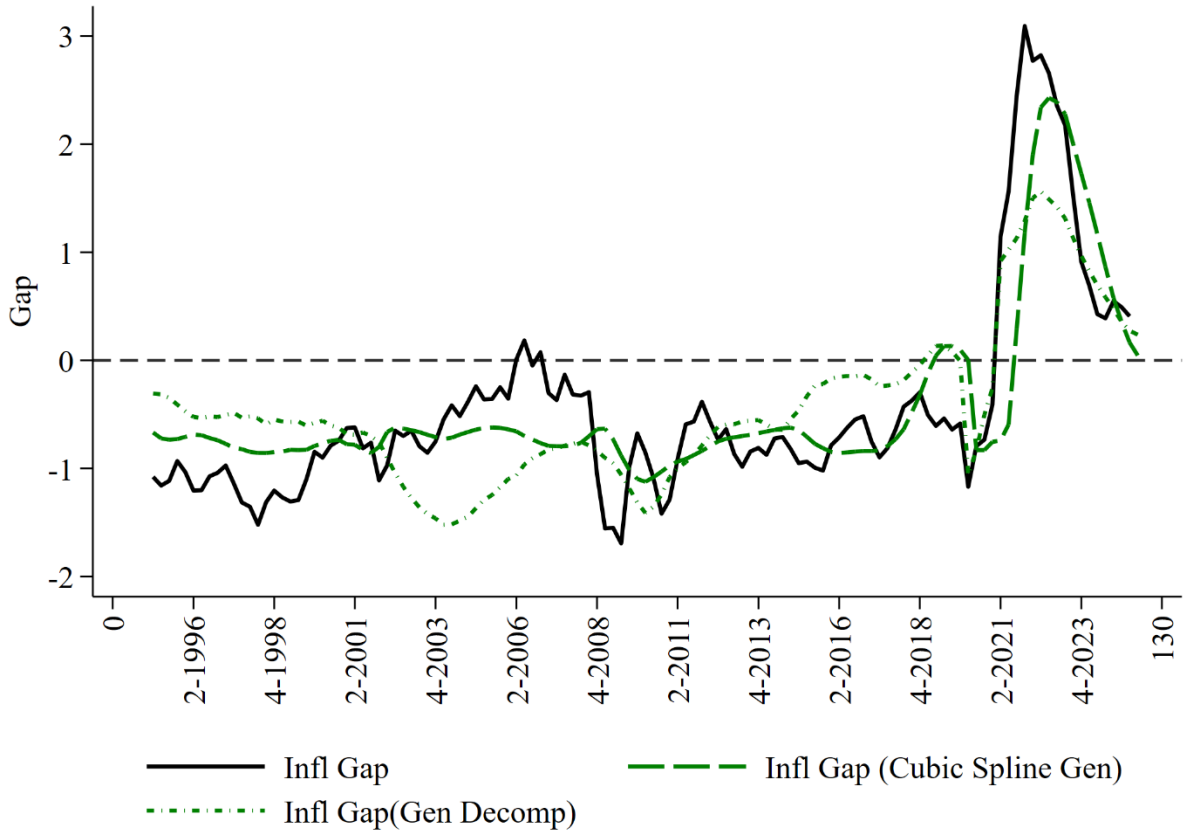


Figure 11: Actual Inflation Gap and Predicted Inflation Gap from Labor Market Tightness: Generalized(Cubic Spline Model) vs. Generalized Decomp.

Notes: Authors' calculations using JOLTS data, CPS data, and the inflation gap measured using PCE core inflation and Livingston survey of inflation expectations. Predicted inflation gap for Cubic Spline Gen from column (4) of Table 1. Predicted inflation gap for Gen Decomp from column (2) of Table 2.

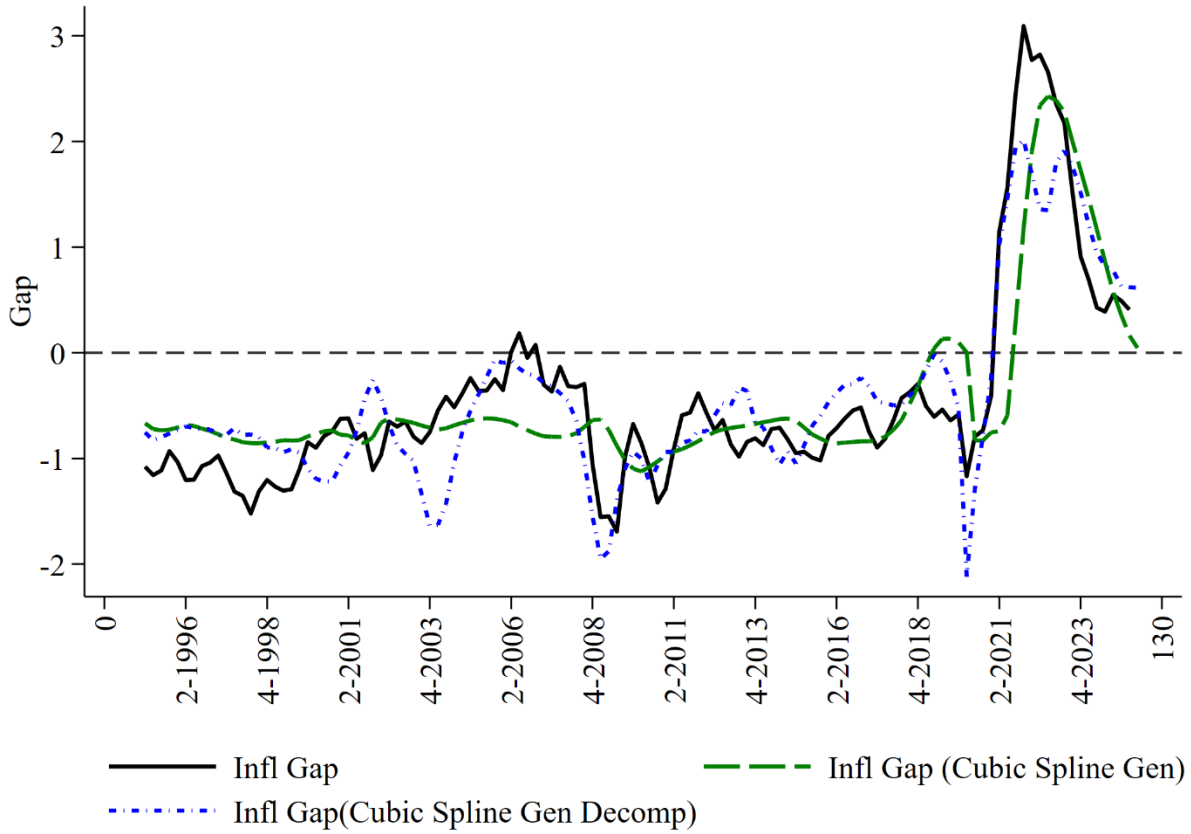


Figure 12: Inflation Gap and Predicted Inflation Gap from (i) Cubic Spline of Generalized Tightness and (ii) Cubic Spline combined with 3-way Decomposition Using Generalized Model

Notes: Authors' calculations using JOLTS data, CPS data, and the inflation gap measured using PCE core inflation and Livingston survey of inflation expectations. Predicted inflation gap from Cubic Spline Gen from column (4) of Table 1. Predicted inflation gap from Cubic Spline Gen Decomp from Table A.2.

Table 1: The Relationship between the Inflation Gap and Labor Market Tightness

	(1) Inflation Gap	(2) Inflation Gap	(3) Inflation Gap	(4) Inflation Gap
Standard Tightness	0.501***			
	(0.187)			
Generalized Tightness			0.583***	
			(0.194)	
Std Tightness (Spline 1)		0.289**		
		(0.138)		
Std Tightness (Spline 2)		-0.984*		
		(0.551)		
Std Tightness (Spline 3)		9.298***		
		(2.932)		
Gen Tight (Spline 1)				0.324**
				(0.138)
Gen Tight (Spline 2)				-1.660**
				(0.670)
Gen Tight (Spline 3)				9.560***
				(2.467)
R-sq.	0.277	0.503	0.375	0.671
Adj. R-sq.	0.271	0.490	0.370	0.663
AIC	297.7	256.0	279.9	205.5
BIC	303.3	267.2	285.5	216.7
Wald χ^2 test	p<0.001		p<0.001	
Obs.	122	122	122	122

Notes: Authors' calculations using JOLTS data, CPS data, and the inflation gap measured using PCE core inflation and Livingston survey of inflation expectations. Newey-West standard errors in parentheses (lags=4). Cubic spline models with 4 interior knots at the 10th, 25th, 65th and 90th percentiles. Wald test for significance of nonlinear spline terms. Constant term not reported. Tightness measures have been z-scored. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2: The Relationship Between the Inflation Gap and Generalized Tightness (using the decomposition into movements along and shifts in the Beveridge curve)

	(1) Inflation Gap	(2) Inflation Gap	(3) Shapley-Owen Contribution from Column (2)
Gen Tightness	0.583***		
	(0.194)		
ES term		-0.133	14%
		(0.148)	
Shifts term		0.863***	21%
		(0.294)	
Matching Eff term		0.962***	61%
		(0.261)	
R-sq.	0.375	0.564	
Adj. R-sq.	0.370	0.553	
AIC	279.9	240.0	
BIC	285.5	251.3	
Wald χ^2 test		p=0.007	
Obs.	122	122	

Notes: Authors' tabulations from JOLTS data, CPS data, and the inflation gap measured using PCE core inflation and Livingston survey of inflation expectations. Column (2) uses decomposition of generalized tightness from equation (16). Newey-West standard errors in parentheses (lags=4). Wald test for equality of coefficients in column (2). Constant term not reported. Tightness measures have been z-scored. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 3: The Relationship between the Inflation Gap (using CPI Core 4Q Inflation) and Tightness (MSA level analysis)

	(1) Inflation Gap	(2) Inflation Gap
Standard Tightness	0.476***	
	(0.093)	
Gen Tightness		0.576***
		(0.094)
R-sq.	0.435	0.444
Adj. R-sq.	0.422	0.431
Obs.	1785	1785

Notes: Authors' calculations using CPS micro data, HWI series from Barnichon and Shapiro (2024), core CPI MSA-level inflation, and Livingston inflation expectations. Standard errors in parentheses clustered at MSA level. Unreported controls are MSA fixed effects, MSA specific trends, relative price of goods to services at the MSA level, NBER recession indicators, Great Financial Crisis recession dummy, pandemic recession dummy, national GDP growth and national core CPI inflation (averaged over 4 prior quarters). Tightness measures are log based averaged over 4 prior quarters and z-score transformed. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix A: Robustness and Sensitivity Exercises

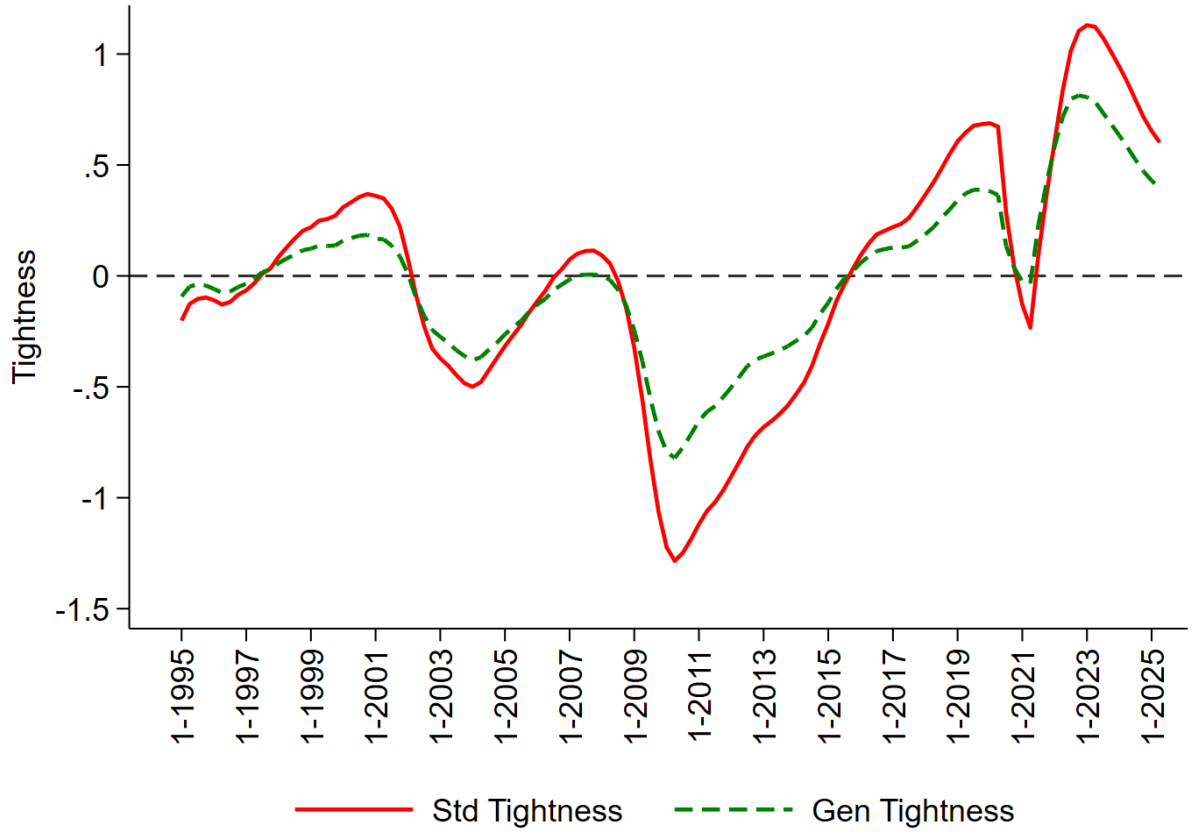


Figure A1: Labor Market Tightness: Standard and Generalized (4 Quarter MA)

Notes: Authors' calculations using JOLTS and CPS data.

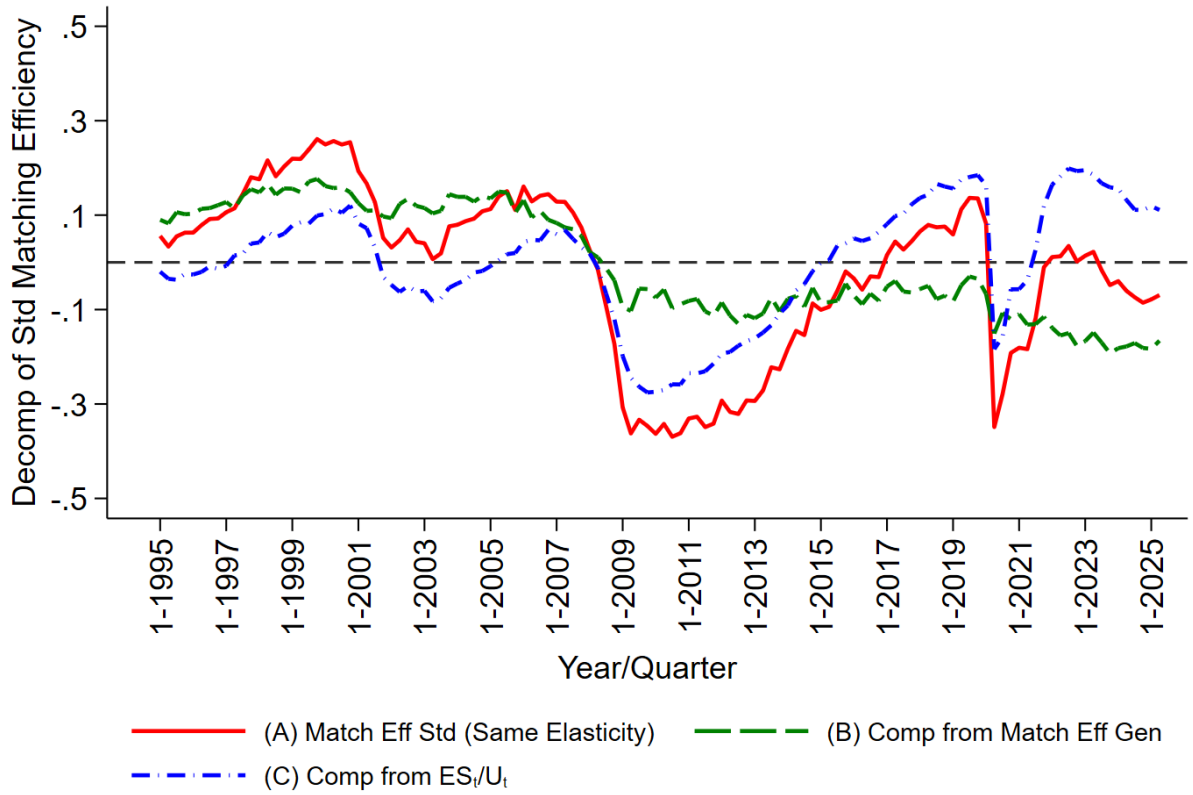


Figure A.2 Robustness Decomposition of Standard Matching Efficiency (Using the Same Elasticity)

Notes: Authors' calculations using JOLTS and CPS data. Decomposition of log matching efficiency from standard model (Match Eff Std) using equation (9) but assuming elasticities are the same at 0.57. Terms are component from generalized matching efficiency (Comp from Match Eff Gen) and component from ES/U (Comp from $\sum_i \rho^{l_i} l_{it} / U_t$).

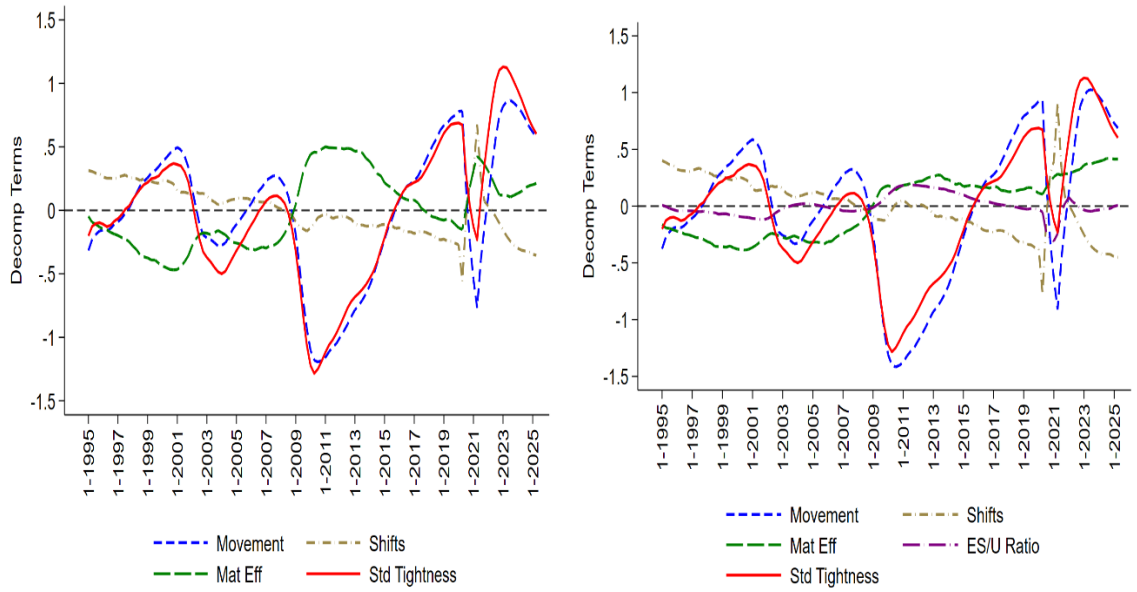


Figure A.3: Decompositions of Traditional Labor Market Tightness: Panel (a) Standard Model; Panel (b) Generalized Model (Using Matching Efficiency from Standard and Generalized Matching Functions)

Notes: Authors' calculations using JOLTS and CPS data.

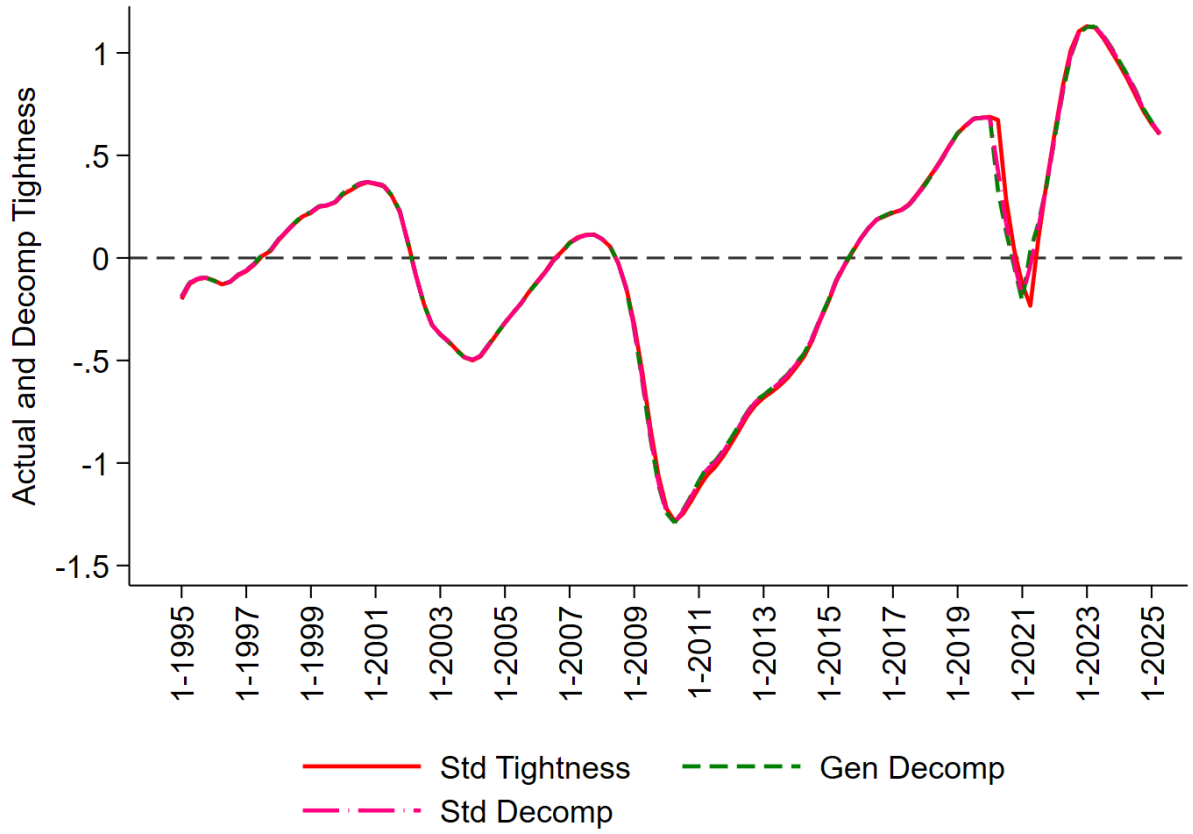


Figure A.4: Decompositions of Standard Labor Market Tightness (Standard and Generalized Approaches Using Matching Efficiency from UE/U flow equations)

Notes: Authors' calculations using JOLTS and CPS data. For Std Decomp, correlation=0.999 and RMSE=0.033. For Gen Decomp, correlation=0.999 and RMSE=0.046.

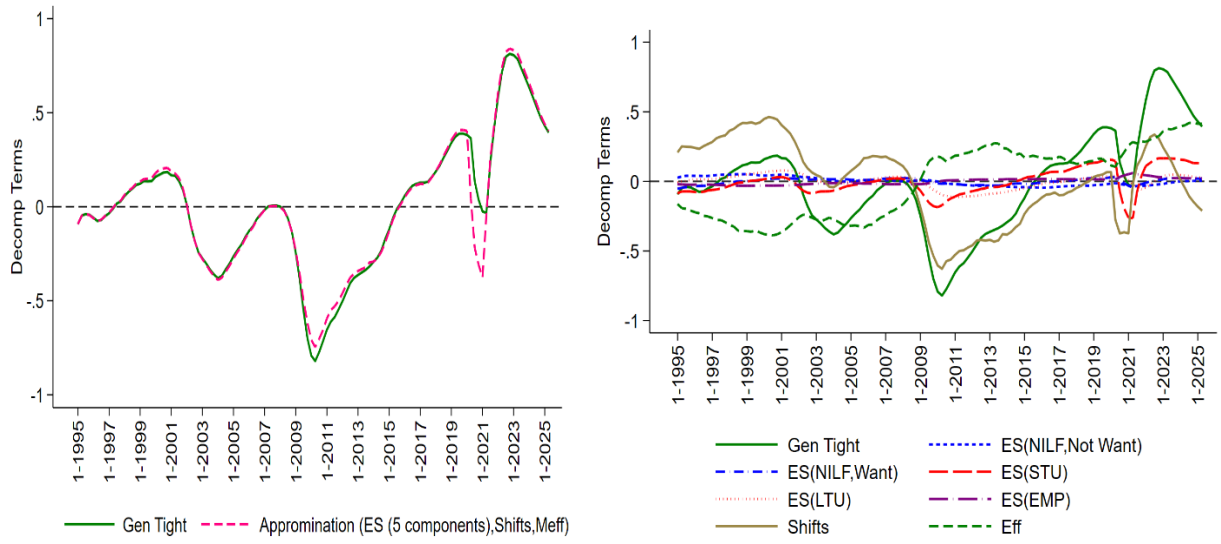


Figure A5: Decompositions of Generalized Tightness (7-way Decomposition)

Notes: Authors' calculations using JOLTS and CPS data. RMSE(Gen-Decomp)=0.066) and Corr(Gen, Decomp)=0.98;

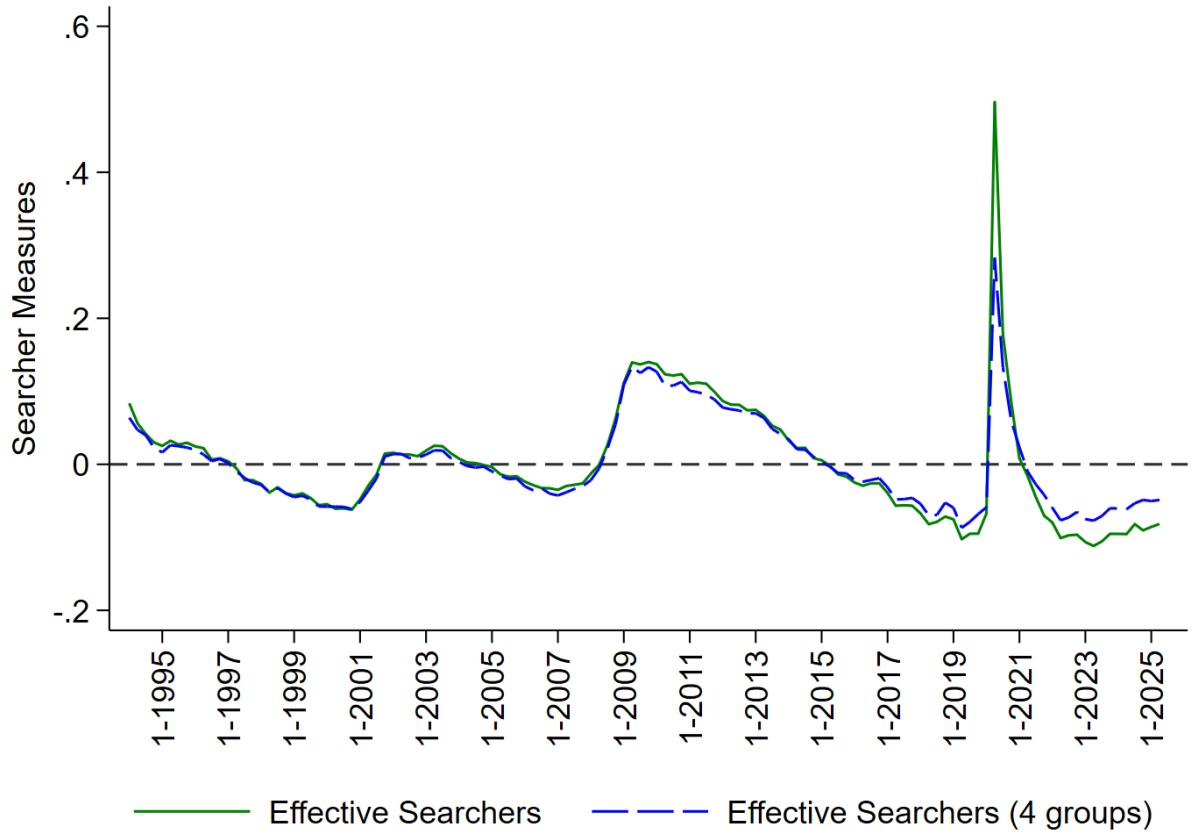


Figure A.6: Comparison of Searcher Measures (Effective Searchers and Effective Searchers (4 groups). Both are rates using working age population. Series reported are logs of rates demeaned.

Notes: Authors' calculations using JOLTS and CPS data.

Table A.1: Estimated Relative Job Finding Rates and Shares in 2006

	Share	Rel. JFR
Unemployed: Recently Left Job	0.16	0.64
Unemployed: Recently Permanently Laid Off	0.22	0.53
Unemployed: Recently Temporarily Laid Off	0.23	1.00
Unemployed: Temp. Job Recently Ended	0.12	0.61
Unemployed: Recently Newly Entered	0.10	0.36
Unemployed: Recently Re-Entered	0.33	0.50
Unemployed: Left Job Months Ago	0.15	0.40
Unemployed: Permanently Laid Off Months Ago	0.37	0.35
Unemployed: Temporarily Laid Off for Months	0.16	0.79
Unemployed: Temp. Job Ended Months Ago	0.14	0.43
Unemployed: Newly Entered Months Ago	0.12	0.36
Unemployed: Re-Entered Months Ago	0.45	0.38
Unemployed: Long-Term Unemployed	0.55	0.29
Want Job: Discouraged	0.17	0.28
Want Job: Looked Last 12 Months	0.45	0.29
Want Job: Other	1.21	0.25
Not in Labor Force: In School	4.41	0.15
Not in Labor Force: Retired	15.23	0.03
Not in Labor Force: Disabled	4.72	0.03
Not in Labor Force: Other	7.46	0.15
Employed: Involuntary Part-Time	1.82	0.08
Employed: Not Involuntary Part-time	61.43	0.03

Notes: The job finding rates are estimated using CPS survey data linking households month-to-month. The relative job finding rate is calculated by dividing all the job finding rates by the job finding rate for the unemployed recently laid off. The recently unemployed groups refer to those who have been unemployed for 0-4 weeks. Unemployed Months Ago refers to those who have been unemployed for 5-26 weeks. Long-Term Unemployed refers to those who have been unemployed for 27 weeks or more.

Table A.2: The Relationship between the Inflation Gap and Labor Market Tightness (Cubic Spline with Decomposition)

ES term (spline 1)	-0.194*
	(0.115)
ES term (spline 2)	1.063**
	(0.418)
ES term (spline 3)	-9.311**
	(3.574)
Shifts term (spline 1)	-0.083
	(0.258)
Shifts term (spline 2)	3.091***
	(0.886)
Shifts term (spline 3)	-16.491***
	(3.490)
Matching Efficiency term (spline 1)	-0.879
	(0.714)
Matching Efficiency term (spline 2)	5.734
	(3.518)
Matching Efficiency term (spline 3)	-6.357
	(4.804)
R-squared (OLS)	0.781
Adj R-squared (OLS)	0.764
ES: Wald χ^2 (s1=s2=s3=0)	.
ES: p-value (all)	0.0043
ES: Wald χ^2 (s2=s3=0)	.
ES: p-value (nonlin)	0.0344
Shifts: Wald χ^2 (s1=s2=s3=0)	.
Shifts: p-value (all)	0.0000
Shifts: Wald χ^2 (s2=s3=0)	.
Shifts: p-value (nonlin)	0.0000
Efficiency: Wald χ^2 (s1=s2=s3=0)	.
Efficiency: p-value (all)	0.0000
Efficiency: Wald χ^2 (s2=s3=0)	.
Efficiency: p-value (nonlin)	0.0000
N	122

Notes: Authors' calculations using JOLTS data, CPS data, and the inflation gap measured using PCE core inflation and Livingston survey of inflation expectations. Newey-West standard errors in parentheses. For each spline, two tests. First, "all" tests whether the joint spline differs from zero. Second, "nonlin" tests whether the nonlinear terms in the decomposition differ from zero.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.3: The Relationship between Inflation Gap and Generalized Labor Market Tightness (Cubic Spline combined with Decomposition)

	Shapley-Owen Contribution from
Decomp of Gen Tightness (Cubic Spline for Each Component)	
ES term in Decomp	20%
Shifts term in Decomp	31%
Matching Eff (Gen) in Decomp	49%
R-sq.	0.781
Adj. R-sq.	0.763
Obs.	122

Notes: Authors' calculations using JOLTS data, CPS data, and the inflation gap measured using PCE core inflation and Livingston survey of inflation expectations.

Appendix B: Median Inflation Robustness Analysis (All results in this section measure the inflation gap=Median Inflation minus Expected Inflation)

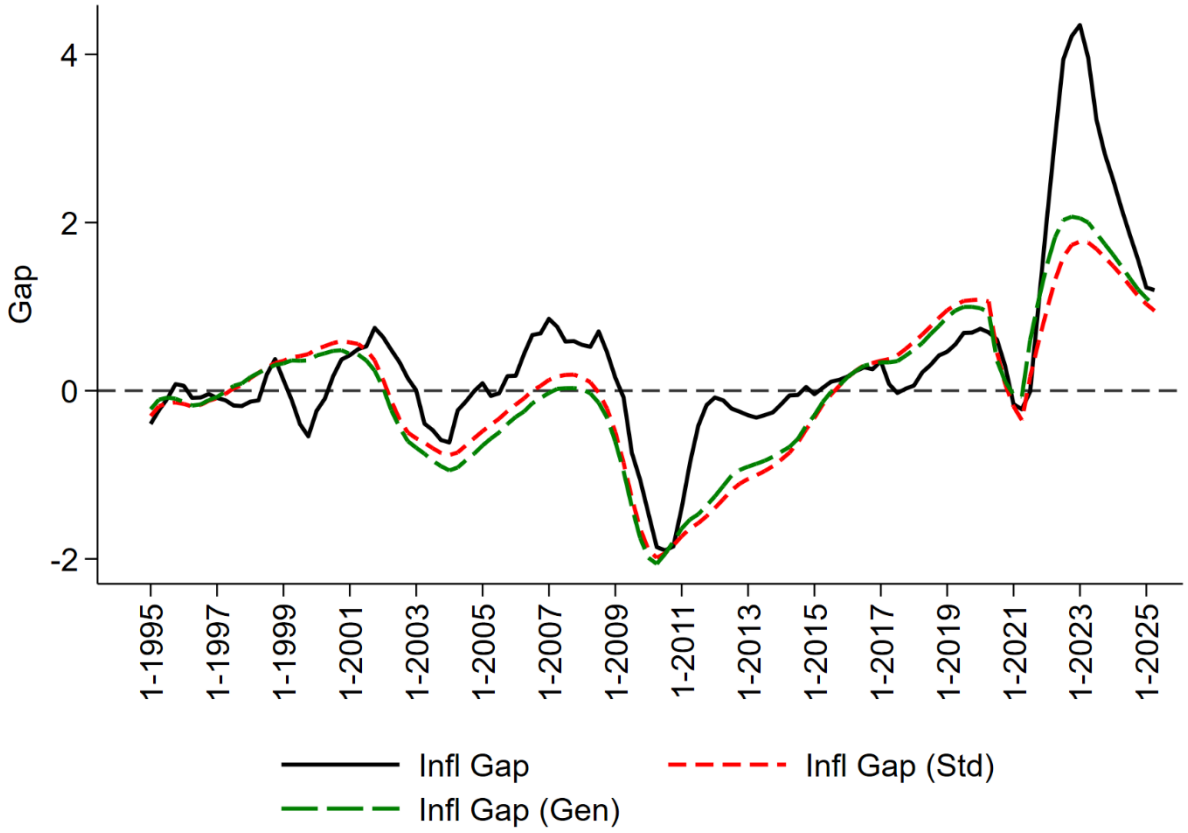


Figure B.1: Actual Inflation Gap and Predicted Inflation Gap from Labor Market Tightness (Standard vs Generalized)

Notes: Authors' calculations using JOLTS national data, CPS micro data and the inflation gap measured using CPI median inflation and inflation expectations from the Livingston survey.

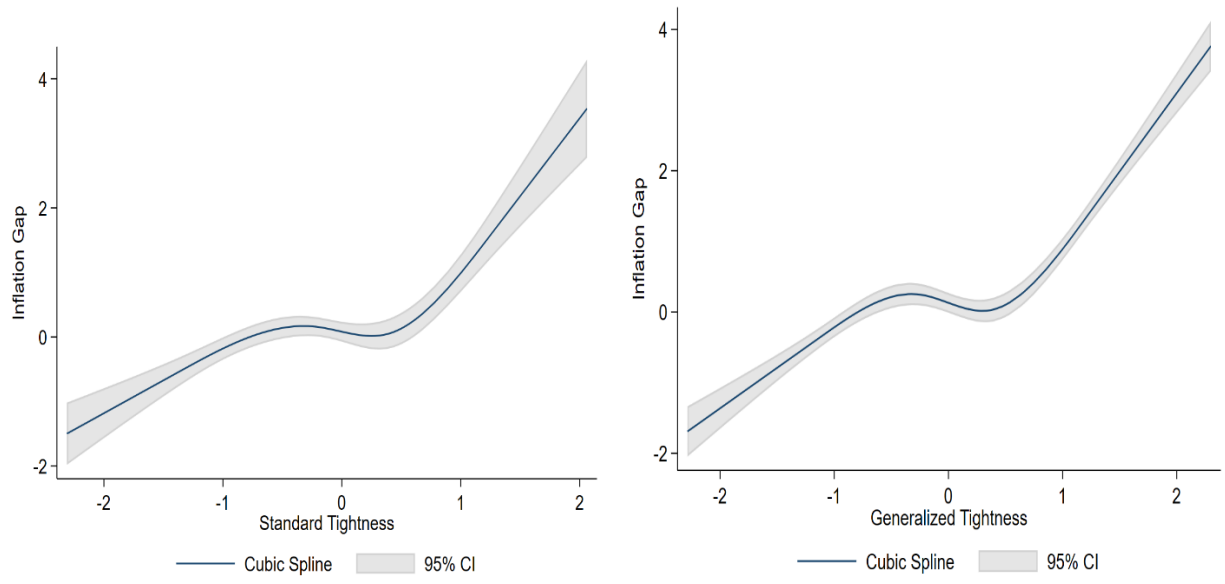


Figure B.2: Nonlinear Relationship between Inflation Gap and Labor Market Tightness (Standard vs. Generalized)

Notes: Authors' calculations using JOLTS national data, CPS micro data and the inflation gap measured using CPI median inflation and inflation expectations from the Livingston survey. Confidence intervals produced using the delta method.

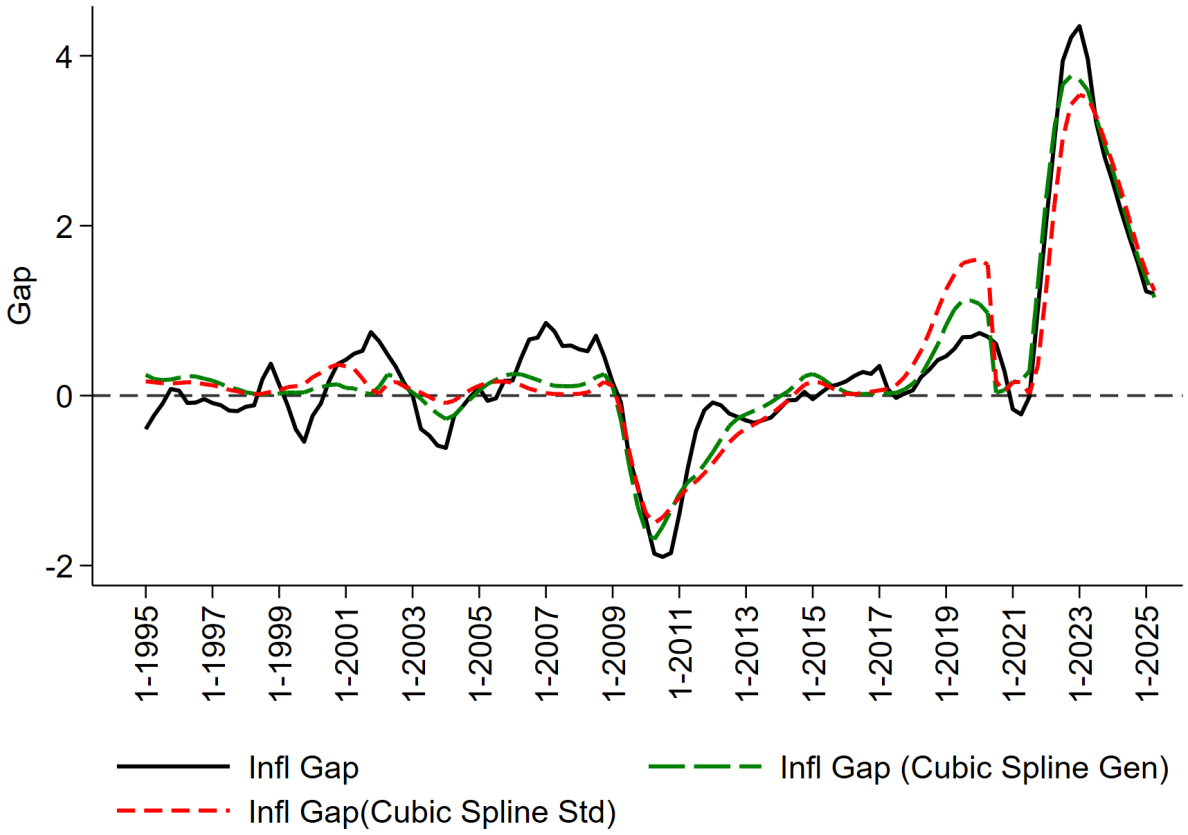


Figure B.3: Actual Inflation Gap and Predicted Inflation Gap from Labor Market Tightness (Standard vs Generalized): Cubic Spline Model

Notes: Authors' calculations using JOLTS national data, CPS micro data and the inflation gap measured using CPI median inflation and inflation expectations from the Livingston survey.

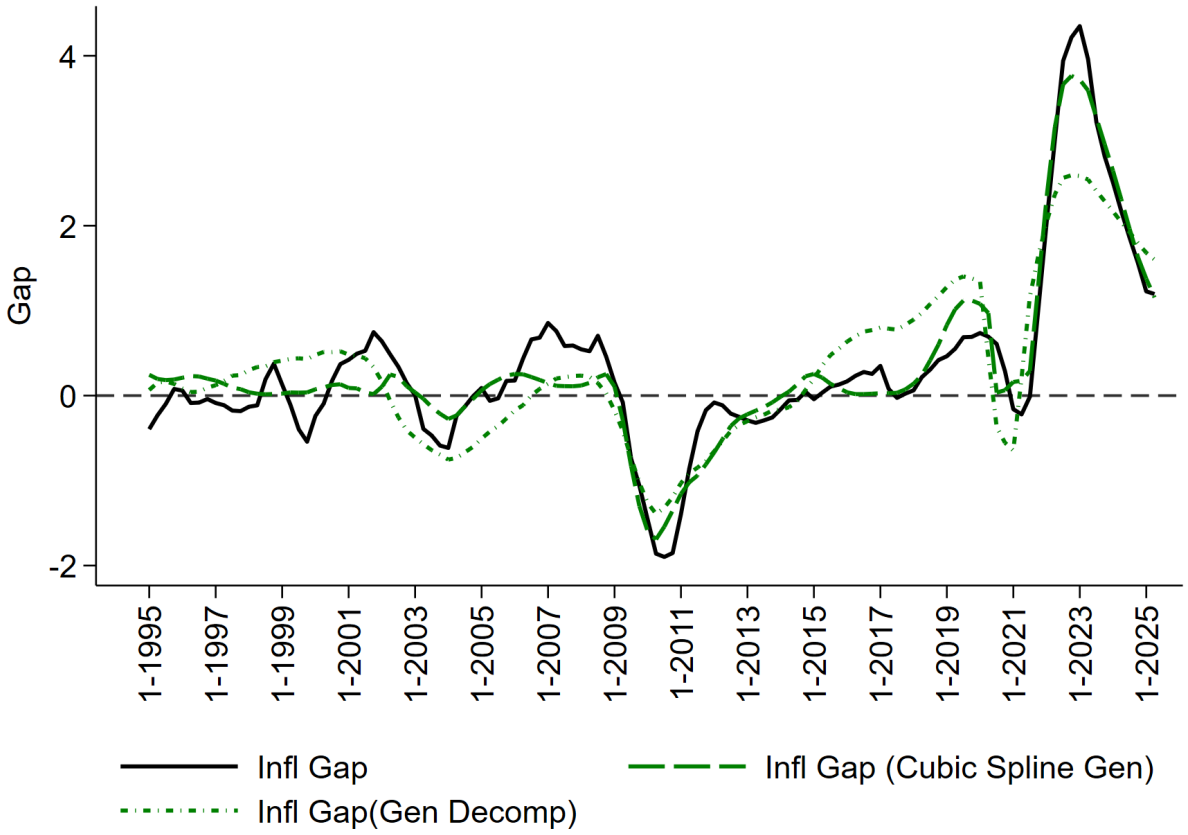


Figure B.4: Actual Inflation Gap and Predicted Inflation Gap from Labor Market Tightness: Generalized(Cubic Spline Model) vs. Generalized Decomp.

Notes: Authors' calculations using JOLTS national data, CPS micro data and the inflation gap measured using CPI median inflation and inflation expectations from the Livingston survey.

Table B.1: The Relationship between the Inflation Gap (Using Median Inflation) and Labor Market Tightness)

	(1)	(2)	(3)	(4)
Standard Tightness	0.860***			
	(0.177)			
Generalized Tightness			0.901***	
			(0.157)	
Std Tightness (Spline 1)		1.019***		
		(0.170)		
Std Tightness (Spline 2)		-2.107***		
		(0.467)		
Std Tightness (Spline 3)		14.653***		
		(2.547)		
Gen Tight (Spline 1)				1.147***
				(0.130)
Gen Tight (Spline 2)				-3.296***
				(0.465)
Gen Tight (Spline 3)				14.742***
				(1.729)
R-sq.	0.649	0.848	0.711	0.914
Adj. R-sq.	0.646	0.844	0.709	0.912
AIC	239.5	141.4	215.7	71.9
BIC	247.9	155.4	224.1	86.0
Wald χ^2 test for nonlinearities		p<0.001		p<0.001
Obs.	122	122	122	122

Notes: Authors' calculations using JOLTS data, CPS data, and the inflation gap measured using CPI median inflation and Livingston survey of inflation expectations. Newey-West standard errors in parentheses (lags=4). Cubic spline models with 4 interior knots at the 10th, 25th, 65th and 90th percentiles. Wald test for significance of nonlinear spline terms. Constant term not reported. Tightness measures have been z-scored. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table B.2: The Relationship Between the Inflation Gap (Using Median Inflation) and Generalized Tightness (using the Decomposition into movements along and shifts)

	(1) Infl Gap	(2) Infl Gap	(3) Shapley-Owen Contribution from (2)
Gen Tightness	0.901***		
	(0.157)		
ES term		0.342***	41.1%
		(0.110)	
Shifts term		0.723***	25.2%
		(0.211)	
Matching Eff		0.845***	33.7%
		(0.203)	
R-sq.	0.711	0.730	
Adj. R-sq.	0.709	0.723	
AIC	213.7	209.7	
BIC	219.3	220.9	
Wald χ^2 test		p=0.061	
Obs.	122	122	

Notes: Authors' tabulations from JOLTS data, CPS data, and the inflation gap measured using CPI median inflation and Livingston survey of inflation expectations. Column (2) uses decomposition of generalized tightness from equation (16). Newey-West standard errors in parentheses (lags=4). Wald test for equality of coefficients in column (2). Constant term not reported. Tightness measures have been z-scored. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Appendix C: Technical Appendix

In this technical appendix we provide more details to derive the log first order approximations of standard and generalized tightness. Combining equations (12) and (13) in the main text, yields:

$$v_t = h_t \left[\frac{\phi_t}{u_t (\mu_t^s)^{1+\eta^s}} \right]^{1/\eta^s} \tag{A1}$$

where $\phi_t \equiv eu_t + nu_t - un_t - \Delta u_{t+1}$. Combining (11) and (13) in the main text yields:

$$v_t = h_t \left[\frac{\phi_t}{u_t (\mu_t^g)^{1+\eta^g} Z_t} \right]^{1/\eta^g} \tag{A2}$$

Combining (4) with $\Delta e_{t+1} = -\Delta u_{t+1} - \Delta n_{t+1}$and $\Delta e_{t+1} = h_t - sep_t$ yields:

$$v_t = \left[\frac{sep_t - \Delta u_{t+1} - \Delta n_{t+1}}{\mu_t^g ES_t^{\alpha_g}} \right]^{\frac{1}{1-\alpha_g}} \tag{A3}$$

Log linearization of (A1) and subtracting $\hat{u}_t$ from both sides (and rearranging terms) yields (14). Log linearization of (A2) and subtracting $\hat{u}_t$ from both sides (and rearranging terms) yields (15). Log linearization of (A3) and subtracting $E\hat{S}_t$ from both sides (and rearranging terms) yields (16).